

**EXPLORING COMMON CRITERIA: CAN IT ASSURE
THAT THE FEDERAL GOVERNMENT GETS NEED-
ED SECURITY IN SOFTWARE?**

HEARING

BEFORE THE

SUBCOMMITTEE ON TECHNOLOGY, INFORMATION
POLICY, INTERGOVERNMENTAL RELATIONS AND
THE CENSUS

OF THE

COMMITTEE ON
GOVERNMENT REFORM

HOUSE OF REPRESENTATIVES

ONE HUNDRED EIGHTH CONGRESS

FIRST SESSION

SEPTEMBER 17, 2003

Serial No. 108-126

Printed for the use of the Committee on Government Reform



Available via the World Wide Web: <http://www.gpo.gov/congress/house>
<http://www.house.gov/reform>

U.S. GOVERNMENT PRINTING OFFICE

92-771 PDF

WASHINGTON : 2004

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov Phone: toll free (866) 512-1800; DC area (202) 512-1800
Fax: (202) 512-2250 Mail: Stop SSOP, Washington, DC 20402-0001

COMMITTEE ON GOVERNMENT REFORM

TOM DAVIS, Virginia, *Chairman*

DAN BURTON, Indiana	HENRY A. WAXMAN, California
CHRISTOPHER SHAYS, Connecticut	TOM LANTOS, California
ILEANA ROS-LEHTINEN, Florida	MAJOR R. OWENS, New York
JOHN M. McHUGH, New York	EDOLPHUS TOWNS, New York
JOHN L. MICA, Florida	PAUL E. KANJORSKI, Pennsylvania
MARK E. SOUDER, Indiana	CAROLYN B. MALONEY, New York
STEVEN C. LATOURETTE, Ohio	ELIJAH E. CUMMINGS, Maryland
DOUG OSE, California	DENNIS J. KUCINICH, Ohio
RON LEWIS, Kentucky	DANNY K. DAVIS, Illinois
JO ANN DAVIS, Virginia	JOHN F. TIERNEY, Massachusetts
TODD RUSSELL PLATTS, Pennsylvania	WM. LACY CLAY, Missouri
CHRIS CANNON, Utah	DIANE E. WATSON, California
ADAM H. PUTNAM, Florida	STEPHEN F. LYNCH, Massachusetts
EDWARD L. SCHROCK, Virginia	CHRIS VAN HOLLEN, Maryland
JOHN J. DUNCAN, JR., Tennessee	LINDA T. SANCHEZ, California
JOHN SULLIVAN, Oklahoma	C.A. "DUTCH" RUPPERSBERGER, Maryland
NATHAN DEAL, Georgia	ELEANOR HOLMES NORTON, District of Columbia
CANDICE S. MILLER, Michigan	JIM COOPER, Tennessee
TIM MURPHY, Pennsylvania	CHRIS BELL, Texas
MICHAEL R. TURNER, Ohio	
JOHN R. CARTER, Texas	
WILLIAM J. JANKLOW, South Dakota	BERNARD SANDERS, Vermont
MARSHA BLACKBURN, Tennessee	(Independent)

PETER SIRH, *Staff Director*

MELISSA WOJCIAK, *Deputy Staff Director*

ROB BORDEN, *Parliamentarian*

TERESA AUSTIN, *Chief Clerk*

PHILIP M. SCHILIRO, *Minority Staff Director*

SUBCOMMITTEE ON TECHNOLOGY, INFORMATION POLICY, INTERGOVERNMENTAL RELATIONS AND THE CENSUS

ADAM H. PUTNAM, Florida, *Chairman*

CANDICE S. MILLER, Michigan	WM. LACY CLAY, Missouri
DOUG OSE, California	DIANE E. WATSON, California
TIM MURPHY, Pennsylvania	STEPHEN F. LYNCH, Massachusetts
MICHAEL R. TURNER, Ohio	

EX OFFICIO

TOM DAVIS, Virginia

HENRY A. WAXMAN, California

BOB DIX, *Staff Director*

CHIP WALKER, *Professional Staff Member*

URSULA WOJCIECHOWSKI, *Clerk*

DAVID McMILLEN, *Minority Professional Staff Member*

CONTENTS

Hearing held on September 17, 2003	Page 1
Statement of:	
Davidson, Mary Ann, chief security officer, Server Technology Platforms, Oracle	73
Fleming, Michael G., Chief, Information Assurance Solutions, Information Assurance Directorate, National Security Agency	21
Gorrie, Robert G., Deputy Director, Defensewide Information Assurance Program Office, Office of the Assistant Secretary of Defense for Net- works and Information Integration, and DOD Chief Information Offi- cer	43
Klaus, Christopher W., chief technology officer, Internet Security Sys- tems, Inc.	83
Roback, Edward, Chief, Computer Security Division, National Institute of Standards and Technology, U.S. Department of Commerce	7
Spafford, Eugene H., professor and director, Center for Education and Research in Information Assurance and Security, Purdue University	88
Thompson, J. David, director, Security Evaluation Laboratory, Cygnacom Solutions	66
Letters, statements, etc., submitted for the record by:	
Clay, Hon. Wm. Lacy, a Representative in Congress from the State of Missouri, prepared statement of	56
Davidson, Mary Ann, chief security officer, Server Technology Platforms, Oracle, prepared statement of	76
Fleming, Michael G., Chief, Information Assurance Solutions, Information Assurance Directorate, National Security Agency, prepared statement of	23
Gorrie, Robert G., Deputy Director, Defensewide Information Assurance Program Office, Office of the Assistant Secretary of Defense for Net- works and Information Integration, and DOD Chief Information Offi- cer, prepared statement of	46
Klaus, Christopher W., chief technology officer, Internet Security Sys- tems, Inc., prepared statement of	85
Putnam, Hon. Adam H., a Representative in Congress from the State of Florida, prepared statement of	4
Roback, Edward, Chief, Computer Security Division, National Institute of Standards and Technology, U.S. Department of Commerce, prepared statement of	10
Spafford, Eugene H., professor and director, Center for Education and Research in Information Assurance and Security, Purdue University, prepared statement of	90
Thompson, J. David, director, Security Evaluation Laboratory, Cygnacom Solutions, prepared statement of	69

EXPLORING COMMON CRITERIA: CAN IT ASSURE THAT THE FEDERAL GOVERNMENT GETS NEEDED SECURITY IN SOFTWARE?

WEDNESDAY, SEPTEMBER 17, 2003

HOUSE OF REPRESENTATIVES,
SUBCOMMITTEE ON TECHNOLOGY, INFORMATION POLICY,
INTERGOVERNMENTAL RELATIONS AND THE CENSUS,
COMMITTEE ON GOVERNMENT REFORM,
Washington, DC.

The subcommittee met, pursuant to notice, at 10:10 a.m., in room 2154, Rayburn House Office Building, Hon. Adam Putnam (chairman of the subcommittee) presiding.

Present: Representatives Putnam, Clay and Watson.

Staff present: Bob Dix, staff director; John Hambel, senior counsel; Chip Walker, professional staff; Ursula Wojciechowski, clerk; Suzanne Lightman, fellow; Erik Glavich, legislative assistant; Ryan Hornbeck, intern; David McMillen, minority professional staff member; and Jean Gosa, minority chief clerk.

Mr. PUTNAM. The Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census will come to order. Good morning, and I apologize for running a few minutes late. I have 20 high school students in Washington for a week for a congressional classroom program to become familiar with the city and our government and how everything works. None of us were figuring on Hurricane Isabel, so we are trying to figure out a way to get 20 airline tickets in very short order, and it's not going to be terribly easy.

Welcome to another important hearing on cybersecurity. Today the subcommittee continues its aggressive oversight and examination of the information security issues most important to our Nation. As many of you know, Secretary Ridge announced the creation of the U.S. Computer Emergency Response Team [U.S.-CERT] in conjunction with Carnegie Mellon University. This is an important step in the progress that needs to be made by our government in protecting the Nation's computers from cyber attack. It's no longer a question of if our computer networks will be attacked, but rather when, how often and to what degree.

Experts from the government and the private sector who have come before this subcommittee are very concerned that the United States is not adequately prepared to ward off a serious cyber attack that could cause severe economic devastation as well as contribute potentially to the loss of life. Blaster and SoBigF are stark examples of how worm and virus vulnerabilities can cost us billions of

dollars in lost productivity and administrative costs in a very short period of time. From the home user to the private enterprise to the Federal Government, we all need to take the cyber threat more seriously and move expeditiously to secure our Nation's computers. I look forward to continuing to work with the Department of Homeland Security and other key Federal agencies in this national security endeavor.

Today's hearing will examine the Common Criteria and whether or not a similar certification should be applied to all government software purchasers. For years countries around the globe have wrestled with the inability to have a commonly recognized method for evaluating security software. Out of this climate, the Common Criteria evolved and represents standards that are broadly useful within the international community.

The international members of the Common Criteria share the following objectives: to ensure that evaluations of information technology products and protection profiles are performed to high and consistent standards, and are seen to contribute significantly to confidence in the security of those products; to improve the availability of evaluated, security-enhanced IT products; to eliminate the burden of duplicating evaluations; and to continuously improve the efficiency and cost-effectiveness of the evaluation and certification/validation process.

The Common Criteria are maintained by an international coalition and is designed to be useful within the widely diverse international community. Currently the recognition arrangement has 15 member countries. The National Security Agency and NIST represent the United States. Each member country accepts certificates issued by the members, making the Common Criteria a global standard. The criteria are technology-neutral and are designed to be applied to a wide variety of technologies and levels of security.

The criteria work by providing standardized language and definitions of IT security components. That standardization allows the consumer, in our case the Department of Defense, to create a customized set of requirements for the security of a product, or protection profile. This profile would include the level of security assurance that the customer desires, including the various mechanisms that must be present for achieving that assurance. Alternatively the criteria allows the producer of the technology to develop their own set of targets called a security target. An independent lab overseen by the participating agencies, in the United States' case NIST and NSA, then test the product against either the profile or the target and certifies that it can satisfy the requirements.

Currently the Department of Defense requires Common Criteria certification for all security-related software purchases. NSA requires Common Criteria certification for all purchases for systems classified as national intelligence.

One of the more useful aspects of the Common Criteria is its ability to allow the purchaser of security software to compare apples to apples. The protection profile which is cast in the language of the Common Criteria provides a view of security features independent of vendor claims. It allows the purchaser to find out with certainty the security features in a product and to compare that

product with other similar ones to determine which ones to purchase.

The certification process, conducted by independent labs overseen by NIST in the United States, concentrates on analyzing the documentation provided by the vendor testing the product, documenting the result and reporting it out to its oversight agency. That agency then reviews the validation report and issues certification. The process is paid for by the vendor and can be both expensive and time-consuming. Estimates for operating systems can be anywhere from 1 to 5 years and costs in the millions of dollars.

The expense and time commitment of the process has given rise to some questioning about the usefulness of the process. For example, the adoption of the Common Criteria could shut small vendors out of the acquisition process because they might not have the resources to go through certification. Another potential problem is the timing. Because certification takes a significant amount of time, the government might not get the most cutting-edge technology available. Conversely, the government does need to gain assurance that security features in products exist and function as advertised.

This is the larger question that we are faced with: How can we—governmentwide—get the most secure products available in a timely and cost-efficient manner and at the same time have IT companies compete on a level playing field in a competitive market that rewards rather than stifles innovation? I look forward to the expert testimony we have assembled today, and I thank the witnesses for their participation.

As with all of our hearings, today's hearing can be viewed live via WebCast by going to reform.house.gov. We will hold off on the other opening statements until the Members arrive, and I would ask that all of our witnesses comply with the light and the timing. Your written statement will be submitted for the record and will be included in its entirety, but we ask that you summarize your verbal comments to 5 minutes.

[The prepared statement of Hon. Adam H. Putnam follows:]

TOM DAVIS, VIRGINIA
 CHAIRMAN
 DAN BURTON, INDIANA
 CHRISTOPHER SMAYS, CONNECTICUT
 L. RAYMOND ROSS, FLORIDA
 JOHN M. MCRAE, NEW YORK
 JOHN L. MICA, FLORIDA
 MARK E. SOUDER, INDIANA
 STEVEN C. LAYCHURTT, OHIO
 DOUG COSE, CALIFORNIA
 RON LEWIS, KENTUCKY
 D. ANN DAVIS, VIRGINIA
 TODD RUSSELL PLATT, PENNSYLVANIA
 CHRIS CANNON, UTAH
 ADAM H. PUTNAM, FLORIDA
 EDWARD L. SCHROCK, VIRGINIA
 JOHN J. DUNCAN, JR., TENNESSEE
 JOHN SULLIVAN, CALIFORNIA
 NATHANIEL, GEORGIA
 CAMDICE MILLER, MICHIGAN
 TIM BURRITT, PENNSYLVANIA
 MICHAEL R. TURNER, OHIO
 JOHN R. CARTER, TEXAS
 WILLIAM J. JANKLOW, SOUTH DAKOTA
 MATTHEW BLACKBURN, TENNESSEE

ONE HUNDRED EIGHTEEN CONGRESS
Congress of the United States
House of Representatives
 COMMITTEE ON GOVERNMENT REFORM
 2157 RAYBURN HOUSE OFFICE BUILDING
 WASHINGTON, DC 20515-6143
 MAJORITY (202) 225-6674
 MINORITY (202) 225-3974
 FAX (202) 225-5051
 E-MAIL (202) 225-6652
 WWW.HOUSE.GOV/REFORM

HENRY A. WAXMAN, CALIFORNIA
 RANKING MINORITY MEMBER
 TOM LANTOS, CALIFORNIA
 MAURICE H. DOWNS, NEW YORK
 EDOLPHUS TOMBS, NEW YORK
 PAUL E. KANARIS, PENNSYLVANIA
 CAROLYN B. MALONEY, NEW YORK
 ELLAH E. CLUMBERG, MARYLAND
 DENNIS J. KUCINICH, OHIO
 DANNY E. DAVIS, ILLINOIS
 JOHN F. TIERNEY, MASSACHUSETTS
 Wm. LACY CLAY, MISSOURI
 DIANE E. WATSON, CALIFORNIA
 STEPHEN F. LYNCH, MASSACHUSETTS
 CHRIS VAN HOLLEN, MARYLAND
 ERIN T. SANDOZ, CALIFORNIA
 C.A. DUTCH BIPPERSBERGER, MARYLAND
 ELEANOR HOLMES MORTON, DISTRICT OF COLUMBIA
 JIM COOPER, TENNESSEE
 CHRIS BELL, TEXAS
 BERNARD SANDERS, VERMONT
 INDEPENDENT

***“Developing Assurance on the Security of Software for the
 Federal Government”***

**Wednesday, September 17, 2003
 10:00 a.m.**

Room 2154 Rayburn House Office Building

Opening Statement of Chairman Adam Putman (R-FL)

Good morning and welcome to another important hearing on cyber security. Today, the Subcommittee continues its aggressive examination of the information security issues most important to our Nation. As many of you know Secretary Ridge announced the creation of the U.S. Computer Emergency Response Team (US-CERT) in conjunction with Carnegie Mellon University.

This is an important step in the progress that needs to be made by our government in protecting the Nation's computers from cyber attack. It is no longer a question of if our computer networks will be attacked, but when, how often, and to what degree. Experts from the government and the private sector who have testified before this Subcommittee are very concerned that the United States is not adequately prepared to ward off a serious cyber attack that could cause severe economic devastation as well as potentially contribute to the loss of life.

Blaster and SoBigF are stark examples of how worm and virus vulnerabilities can cost us billions of dollars in lost productivity and administrative costs in a very short period of time. From the home user, to private enterprise, to the Federal government, we all need to take the cyber threat more seriously and move expeditiously to secure our Nation's computers. I look forward to continuing to work with DHS and other key federal agencies such as OMB, DOD, NIST, and NSA in this national security endeavor.

Today's hearing will examine the Common Criteria and whether or not a similar certification should be applied to all government software purchases. For years countries around the globe have wrestled with the inability to have a commonly recognized method of evaluating security software. Out of this climate the Common Criteria evolved and represents standards that are broadly useful within the international community

The international members of the Common Criteria share the following objectives:

1. to ensure that *evaluations of Information Technology (IT) products and protection profiles* are performed to high and consistent standards and are seen to contribute significantly to confidence in the security of those products and profiles;
2. to improve the availability of evaluated, security-enhanced IT products and protection profiles;
3. to eliminate the burden of duplicating evaluations of IT products and protection profiles;
4. to continuously improve the efficiency and cost-effectiveness of the evaluation and *certification/validation** process for IT products and protection profiles.

The Common Criteria are maintained by an international coalition and is designed to be useful within the widely diverse international community. Currently, the Common Criteria Recognition Arrangement has 15 member countries. The National Security Agency and the National Institute for Standards and Technology represent the U.S.

Each member country accepts certificates issued by other members, making the Common Criteria a global standard. The criteria are technology neutral and are designed to be applied to a wide variety of technologies and levels of security.

The Criteria work by providing standardized language and definitions of IT security components. That standardization allows the consumer, in our case, the Department of Defense to create a customized set of requirements for the security of a product (called a protection profile). This profile would include the level of security assurance that the customer desires, including the various mechanisms that must be present for achieving that assurance. Alternatively, the Criteria allows the producer of the technology to develop their own set of targets (called a security target). An independent lab, overseen by the participating agencies (NIST and NSA in the U.S.) then tests the product against either the profile or the target and certifies that it can satisfy the requirements. Currently, the Department of Defense requires Common Criteria Certification for all security-related software purchases. NSA requires Common Criteria certification for all purchases for systems classified as national intelligence.

One of the more useful aspects of the Common Criteria is its ability to allow the purchaser of security software to compare "apples to apples." The protection profile, which is cast in the language of the Common Criteria, provides a view of security features independent of vendor claims. It allows a purchaser to find out, with certainty, the security features in a product, and to compare that product with other similar ones to determine which one to purchase.

The certification process, conducted by independent labs overseen by Common Criteria members (NIST in the U.S.), concentrates on analyzing the documentation provided by the vendor, testing the product, documenting its result and reporting out to its oversight agency. The oversight agency then reviews the validation report and issues a certification. The certification process is paid for by the vendor and can be both expensive and time consuming. Estimates for operating systems can be anywhere from 1-5 years and cost millions of dollars.

The expense and time commitment of the process has given rise to some questioning of the usefulness of the process. For example, the adoption of Common criteria could shut small vendors out of the acquisition process because they might not have the resources to go through certification. Another potential problem is timing. Because certification takes a significant amount of time, the government might not get the most cutting-edge technology available. Conversely, the government does need to gain assurance that security features in products exist and function as advertised.

This is the larger question that we are faced with: How can we -- government-wide -- get the most secure products available in a timely and cost efficient manner and at the same time have IT companies compete on a level playing field in a competitive market that rewards and doesn't stifle innovation?

I look forward to the expert testimony we will hear today and thank the witnesses for their participation.

Today's hearing can be viewed live via WebCast by going to **<http://reform.house.gov>** and then clicking on the link under "Live Committee Broadcast".

Mr. PUTNAM. With that, as is the custom of this subcommittee, we will swear in the witnesses. I will ask our first panel rise and raise your right hands.

[Witnesses sworn.]

Mr. PUTNAM. Note for the record all of the witnesses responded in the affirmative. And we will move right to our distinguished panel.

Our first witness is Edward Roback. Mr. Roback serves as the Chief of the Computer Security Division at the National Institute of Standards and Technology supporting the agency's responsibility to protect sensitive Federal information and promote security and commercial information technology products. As Chief, he leads the implementation of NIST responsibilities under FISMA and Cybersecurity Research and Development Act. Mr. Roback heads NIST's participation on the NIST-NSA Technical Working Group and serves on the Committee of National Security Systems. He has chaired the Federal Agency Computer Security Programs Managers Forum and co-authored *An Introduction to Computer Security, The NIST handbook*. He has also recently authored NIST's *Guidelines to Federal Organizations on Security Assurance and Acquisition/Use of Tested/Evaluated Products*. For those of who you would like a copy, they will be available at Barnes and Noble afterwards, and he will be happy to autograph them for you.

Mr. Roback, you are recognized for 5 minutes. Welcome to the subcommittee.

STATEMENT OF EDWARD ROBACK, CHIEF, COMPUTER SECURITY DIVISION, NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, U.S. DEPARTMENT OF COMMERCE

Mr. ROBACK. Thank you, Chairman Putnam. Thank you for the opportunity to testify today. In response to your invitation, I first would like to discuss what security assurance is and the role it plays in overall cybersecurity. I then would like to turn to the role that security testing and particularly the Common Criteria and the NIST-NSA-NIAP program play. I would like to leave you with some ideas as to what else the research community can do to improve the trust and confidence that we must have in the proper, correct and secure functioning of information systems. So let me start.

What is security assurance? If we look at assurance broadly, it's the basis we need for overall trust and confidence in the correct and secure information systems. The overall question of assurance tries to address two questions: Does the system do what it is supposed to do, and does it not do the unintended? Within that context, security assurance, simply put, is the degree of confidence one has that the security mechanisms of a system is intended. It is not an absolute guarantee that security is achieved.

How do we get security assurance? There is no single way. One can get some degree by looking at how a system is built, the past use of a system, manufacturers' warranties or lack thereof, and, of course, independent testing and evaluation. This testing can vary from the straightforward and repeatable through the more complex and time-consuming. When we have a standard specification that is very precise, such as with an encryption algorithm, testing is

straightforward, although not necessarily easy. When a specification is exact, the test can be correspondingly precise.

On the other hand, when we look at more complex and diverse IT products which lack common standards specification at the bits and bytes level, we're often confronted with products containing millions of lines of code for which a standard spec does not exist, and testing is not just straightforward. Testing such products necessarily involves human subjectivity. NIST refers to such testing as evaluation. NIAP is such a testing program.

Turning to the NIAP and the CC, in my written statement I have provided a summary of the development of each, the Mutual Recognition Arrangement and some of the uses of the criteria both domestically and overseas, and indeed there have been some very significant uses. Major issuers of bank cards have formed work groups to use the Common Criteria to develop a profile for smart cards; the Financial Services Business Roundtable is doing that for the financial services community. The Process Control Security Requirements Forum is using the Common Criteria for SCADA systems security, and it is also being used in the health care community trying to use the Common Criteria to define requirements for the health care systems.

But I think it's important to take a minute to review the meaning of a Common Criteria certificate. A Common Criteria evaluation is a measure of the information technology's compliance to the vendor's claimed security. It is not a measure or a guarantee that the product is free from malicious code or that the overall comprised system is secure. Any product that has a Common Criteria specification can undergo an evaluation and receive its certificate if the evaluation process is completed. I provided additional details in my written statement.

As you mentioned, we have issued advice to the agencies on the use of evaluated products for non-national security systems. We described the overall role that assurance can play. And, of course, the Committee for National Security Systems has issued its Policy No. 11, and I will defer to my colleagues for additional comments on that.

As to whether that policy should be extended, I believe that more data is needed from the CNSS policy experience before extension is considered or recommended for unclassified systems. One of the criticisms often levied on NIAP is that evaluations take too long and cost too much. We hear this from the small business community. Of course, one would expect to hear that of any evaluation process that is not free and instantaneous. However, these products do involve millions of lines of code. But given resolve, flexibility, resources and research, significant progress can be made.

For example, the research community should look at new ways to develop enhanced security testing. We need new methods. The current process we have is too expensive and involves too much human subjectivity. We need to invest more in doing such research, because the sooner we do, the sooner we will have benefits from the results. We need to look outward at system-level composability issues and enterprise architecture issues, and we need to look inward to some of the security issues that are present with things like protocols. You have to look across the entire spectrum.

In summary, the Common Criteria provides the means to develop specifications and a common means to develop security evaluations. However, more can be done to streamline this process through research and standards development, resources permitting. We must also keep in mind that technology alone will not achieve security, although we are focused on technology today.

Thank you for the opportunity to testify today.

Mr. PUTNAM. Thank you very much.

[The prepared statement of Mr. Roback follows:]

Statement of

Edward Roback

**Chief, Computer Security Division
National Institute of Standards and Technology**

U.S. Department of Commerce

Before the

**Committee on Government Reform
Subcommittee on Technology, Information Policy,
Intergovernmental Relations and the Census**

**“Exploring Common Criteria: Can it Ensure that the Federal
Government Gets Needed Security in Software?”**

September 17, 2003

Chairman Putnam, Representative Clay, and Members of the Subcommittee, thank you for this opportunity to testify today. The Computer Security Division at the National Institute of Standards (NIST) has direct responsibility for NIST's activities associated with Common Criteria and the National Information Assurance Partnership. In response to the issues raised in the letter of invitation, I would like to first discuss what security assurance is and the role it plays in overall cyber security. I then will turn to the role that security testing, and specifically the Common Criteria (CC) and the NIST-National Security Agency (NSA) National Information Assurance Partnership (NIAP), play in helping to bring about security assurance. Finally, I would like to leave with you some ideas as to what else the cyber security research community could do to improve the trust and confidence we have in the proper, correct, and secure functioning of information systems.

Security Assurance

Assurance is the basis we need for overall trust and confidence in the correct *and* secure operation of information systems. The overall question of assurance tries to address two important questions: Does a system do what it is supposed to do? And, does the system do anything that is unintended? Within this context, *security assurance*, simply put, is the degree of confidence one has that the security measures of a system work as intended; it is *not* an absolute guarantee that security is achieved. We need to keep this in mind when discussing the NIAP, or any other security testing program. Today I will be speaking primarily to the question of security assurance, within this overall context.

Why is security assurance important? The risks we decide to take with regard to systems are based upon the system vulnerabilities and an assessment of potential losses if such vulnerabilities become manifest. (There are formal definitions of "risk levels" in the security community, but I am using the term in a more general sense here.) This can be clearly seen with life-critical systems. We generally are not willing to accept the potential losses from failure of a life-critical system! Rather, a high degree of confidence is required in the correct and secure operation of a system that could result in a loss of life. If we have good reasons to be confident in the security of a system, we can reasonably be expected to rely upon the system for more important tasks and the processing of more sensitive information. In the Federal context, security assurance is an important input to the security accreditation process, namely the decision by a management official to place a system into operation.

How is security assurance obtained? There is no single way. One can gain *some degree* of confidence in the security of a system (or component, etc.) by looking at the process of how the system is built. If a rigorous methodology of requirements definition, design specification, and conformance or acceptance testing is in place, one would generally have more confidence in the resulting system than one developed haphazardly. Similarly, use of advanced software engineering techniques can provide assurance. The past experience of use of a particular system is another means by which one can gain *some degree* of assurance. If a system is used by a hundred organizations without security incidents (which, by the way, can be most difficult to ascertain), one can make a

reasonable leap-of-faith that it will also operate securely in the hundred-and-first. Manufacturers' warranties or lack thereof is another means to have *some degree* of security assurance. Ensuring the continued security of a system once in operation is also important. Scanning tools can be (and should be) used to help ensure that important security settings are maintained and that known vulnerabilities are located and patched. There are many other means as well to help obtain and maintain security assurance. Of course, last but not least, is the use of independent security testing and evaluation to help achieve security assurance.

Security Testing and Evaluation

Security testing can be achieved through a range of means from the straightforward and repeatable through more complex and time consuming processes.

When a standard specification exists, such as an encryption algorithm, it is a reasonably straightforward (but not necessarily easy) process to determine whether the algorithm is correctly implemented. In this case, the specification is exact, and the tests can be correspondingly precise. NIST refers to this process as conformance testing and *validation*. I should note here that the Cryptographic Module Validation Program operated by NIST and the Communications Security Establishment of the Government of Canada provides such algorithm and related testing.

On the other hand, as we look at more complex and diverse information technology (IT) products lacking common/standard specifications, we are often confronted with products containing millions of lines of software code for which a standard bits-and-bytes level specification does not exist. Testing such products necessarily involves human subjectivity; NIST refers to such testing as *evaluation*. That is not to say evaluation cannot be and is not rigorous; it certainly can and probably should be more rigorous than current practices (depending upon the level of effort and time one wishes to expend.) What I am saying is that such testing is considerably removed from more straightforward, "black-box", yes/no testing. Although there is promise for the use of formal methods here, today the use of such techniques is considered by vendors to be expensive. Formal methods are of particular note as they can both be used to increase the quality of software and to facilitate the automatic generation of tests, including expected outputs, from formal specifications. A 2002 NIST commissioned study of the economic impact of software quality showed that software bugs, or errors, are so prevalent and so detrimental that they cost the U.S. economy an estimated \$59.5 billion annually, or about 0.6 percent of the gross domestic product. Findings of the 309-page report are intended to identify the infrastructure needs that NIST can meet through its research programs. Though assurance programs can be built by various sectors NIST's programs address assurance, trust and confidence in general.

Next, let me turn more specifically to the NIST-NSA NIAP program, which provides security evaluation of IT products and is built upon the use of the CC.

Common Criteria

Development of the CC began in 1993 in response to efforts by a range of nations to develop IT security evaluation criteria. Efforts were underway in Canada, the U.K. and the E.U. to develop such criteria at the same time the US was considering a revision to the 1985 Department of Defense evaluation criteria commonly known as the "Orange Book." The development of different sets of criteria, which were not harmonized, presented costly potential conflicts to the IT industry. Vendors were going to be faced with the need to undergo multiple security evaluations in multiple countries. The likelihood of non-tariff barriers to trade loomed large. For this reason, security experts from NIST and NSA partnered with the U.K., Canada, Germany, France and the Netherlands and set a goal of developing a single set of criteria under which security evaluations could take place.

In May of 1998, the CC was completed. The 800-plus page document is known formally as ISO/IEC 15408: Common Criteria for Information Technology Security Evaluation. It is intended for use for either the specification of security requirements (i.e., properties) of a product (e.g., a specification for the security capabilities in a firewall), as the basis for security evaluation of security requirements of IT products and systems, or both.

As a security requirements specification language, the CC enables user communities (e.g. health care, financial, SCADA) to state to technology providers what security capabilities they desire in products they wish to buy. In addition, developers of specific products can use the CC to tell potential customers exactly what security capabilities are contained in the product.

As the basis for the evaluation of security requirements, the CC permits comparability between the results of independent security evaluations. It does so by providing a common taxonomy of security functional requirements for describing IT products and systems and of assurance measures that are applied during development and evaluation of the products/systems. The evaluation process establishes a level of confidence that the products and systems conform to their stated security functional and assurance requirements, which have been specified using the CC. The evaluation results are intended to help consumers determine whether the IT product is secure enough for their intended application and whether the security risks are acceptable.

The great potential of the CC is both in (1) its use to express "good sets of requirements" and (2) to provide assurance, through evaluation, that products comply with these requirements. Examples of how various user communities have and are using the CC to state its security requirements are given later. Unfortunately, the use of the CC as a requirements specification language has been under-utilized.

Common Criteria Mutual Recognition Arrangement

The completion of the CC was followed by the signing of the CC Recognition Arrangement (CCRA), now including 17 signatory nations, in order to reduce the cost of multiple evaluations to vendors. In October 1998, Government organizations from the United States, Canada, France, Germany, Netherlands, and the United Kingdom signed an historic mutual recognition arrangement for Common Criteria-based evaluations. The Arrangement, officially known as the *Arrangement on the Recognition of Common Criteria Certificates in the field of Information Security*, was a significant step forward for Government and industry in the area of IT product security evaluations. The partners in the Arrangement share the following objectives in the area of Common Criteria-based evaluations of IT products:

- To ensure that evaluations of IT products are performed to high and consistent standards and are seen to contribute significantly to confidence in the security of those products;
- To increase the availability of evaluated, security-enhanced IT products for national use;
- To eliminate the need for redundant evaluations of IT products; and
- To continuously improve the efficiency and cost-effectiveness of security evaluations and the validation process for IT products.

The purpose of this Arrangement is to advance those objectives by bringing about a situation in which security-enhanced IT products that earn a Common Criteria certificate can be procured or used without the need for them to be evaluated and validated again. It seeks to provide grounds for confidence in the reliability of the judgments on which the original certificate was based by declaring that the Validation Body associated with a Participant to the Arrangement shall meet high and consistent standards. The Arrangement specifies the conditions by which each Participant will accept or recognize results of IT security evaluations and the associated validations conducted by other Participants and to provide for other related cooperative activities.

Since its original signing, Australia, New Zealand, Greece, Finland, Israel, Italy, Spain, Norway, Austria and Sweden have signed the arrangement. In addition, a number of countries such as Japan, Russia, and Korea have indicated their intent to accede to the arrangement.

National Information Assurance Partnership

As the CC was nearing completion, NIAP was created in 1997 by NIST and NSA to bring together the technical expertise from both agencies to focus on the development of cost-effective testing and evaluation techniques and methods for assessing the security features in commercial off-the-shelf IT products. The partnership emphasized the use of the CC, the involvement of other industrialized nations beyond the United States in recognizing the results of the security evaluations performed, and the participation of private industry, whenever possible, in developing security-enhanced IT products and in

conducting security evaluations. In the U.S., NIAP security evaluations are conducted by commercial testing laboratories that have been accredited under NIST's National Voluntary Laboratory Accreditation Program.

The NIAP Validation Body assesses the results of a security evaluation conducted by a testing lab and issues a CC certificate. The certificate, together with its associated validation report, confirms that an IT product has been evaluated at an accredited testing laboratory using the Common Methodology for conformance to the CC. The certificate also confirms that the IT security evaluation has been conducted in accordance with the provisions of the testing program and that the conclusions of the testing laboratory are consistent with the evidence presented during the evaluation that the product conforms to its security specification. I should note, the certificate does not mean that the product is necessarily secure. I will speak more about that later.

NIAP maintains a Validated Products List on its web site containing all IT products that have successfully completed evaluation and validation under the testing program. The validated products list also includes those products that have successfully completed similar processes under the testing programs of authorized signatories to the CC MRA.

Today, NSA leads the day-to-day operations of the Validation Body, that is, NSA reviews and validates the test results and issues the CC certificate for the vendor's product based on the lab assessment. NIST leads the laboratory accreditation program bringing in new laboratories to the testing program and re-accrediting the current network of CC testing labs. Given resource constraints, this division of labor and responsibilities for the testing program seems to be the most effective method of allocating resources.

The Meaning of a NIAP (or Other) Common Criteria Certificate

As I mentioned earlier, it is important to understand exactly what CC evaluation, and specifically a CC certificate means. A CC evaluation is a measure of an information technology product's compliance to the vendor's claimed security (specification using the Common Criteria). It is not a measure of how much protection the claimed security specification provides nor does it guarantee that the product is free from malicious or erroneous code. Any product that has a CC security specification can undergo an evaluation and receive a certificate if it successfully completes the evaluation. It is important for users to understand what the issuance of a CC certificate does and does not imply. A CC certificate:

- **Does** mean that NIST and NSA (or equivalent government organizations participating in the CCMRA) believe the evaluation has been conducted properly and the conclusions of the private sector testing laboratories are consistent with the evidence produced.
- **Does** imply that a good faith effort has been made to ensure that the product conforms to the security claims stated by the vendor in the security specification.
- **Does not** imply **with absolute certainty** that the product conforms to the security claims stated by the vendor in the security specification.

- **Does not** imply that the product conforms to security claims in documents other than the security specification (i.e., security claims in promotional literature, vendor documentation, and other documents **are not** covered by the validation certificate).
- **Is not** an endorsement or warranty of the product by NSA and NIST (or by equivalent government organizations participating in the CCRA).
- **Does not** imply or guarantee that the product is free from malicious or erroneous code.
- **Does not** imply that security functional specifications and achieved level of assurance of the product provide adequate protection for data contained in the product's intended operational environment.
- **Does not** presume that subsequent versions or releases of the product should not be or do not have to be evaluated.

Upon successful completion of a CC evaluation, the product's security specification and the Validation Report are posted to the NIAP website (<http://niap.nist.gov/cc-scheme/ValidatedProducts.html>) to allow consumers to confidently make acquisition decisions regarding different products.

Use of the Common Criteria

Within the U.S. Federal Government, the use of CC and NIAP- evaluated products is addressed by NIST through its advice to agencies for non-national security systems through "Guideline to Federal Organizations on Security Assurance and Acquisition/Use of Tested/Evaluated Products," (See NIST Special Publication 800-23, available at <http://csrc.nist.gov/publications/nistpubs/index.html>). This publication describes how assurance in acquired products supports security and the benefits that can be obtained through testing of commercial products against customer, government, or vendor-developed specifications. Also discussed is the need for Federal departments and agencies to acquire and use products appropriate to their risk environment while considering cost-effective selection of security measures. NIST recommends that Federal agencies give substantial consideration in IT procurement and deployment for IT products that have been evaluated and tested by independent accredited laboratories against appropriate security specifications and requirements. The Committee for National Security Systems (CNSS) has issued its CNSS Policy #11, recently amended, to address national security systems, and I will defer to my colleagues from that community to address it. The potential extension of CNSS Policy #11 beyond the national security community may be addressed as part of the national review of NIAP called for in the White House's *National Strategy to Secure Cyberspace* (February 2003). However, more data is needed on the impact of the policy before extension is considered or recommended. As the national security community gains experience from its policy, one can consider whether it should be extended to non-national security systems.

Other governments are also adopting, on either a voluntary or regulatory basis, the use of the CC. France has in place a regulation recommending use of CC evaluations for public administration. The European Union has passed a resolution on information and network

security addressing use of the CC for electronic signatures. The CC has been adopted by NATO as a standard. In Germany CC evaluations are required in their digital signature legislation.

Use of the CC by User Communities to state their security requirements

As mentioned earlier, we believe the most under-utilized aspect of the CC is as a requirements specification language. While there are some excellent examples of such use, the full benefits of the CC will not be achieved until there is a better balance between its use for evaluation and for security requirements specification. When used as requirements specification language, the CC allows communities-of-interest that procure IT products to state the security requirements they wish to have developers supply in products. The security requirements can be for technology-specific products or for application-oriented use. As an example of technology specific security requirements, NIST and NSA are developing security requirements for technologies such as firewalls, intrusion detection systems, biometrics, and operating systems. The security requirements are developed using the CC Protection Profile construct. These profiles are statements by NIST and NSA about what “good” security requirements are for these technologies.

As examples of application-oriented Protection Profiles, we cite:

- The major bankcard issuers (e.g., American Express, Mastercard, Visa) formed a working group that used the CC to develop a profile for the smartcards they issue to their customer banks. A significant effort (the first of this type) was the group’s development of their profile for smartcards.
- The Financial Services Roundtable/BITS, whose members consist of major banks and insurance companies, has used the CC to specify the security functionality its members would like to see in various IT products. When a product that meets BITS security functionality receives a CC certificate, BITS will issue its mark on that product based on the CC evaluation that was performed.
- The Process Control Security Requirements Forum (PCSRF), led by NIST, is composed of government and private sector representatives who are defining security requirements for products used in real-time processing and SCADA systems. The goal of this effort is to influence the key vendors that supply products and systems globally for real-time and SCADA systems to meet process control security requirements. If vendors respond to these market signals, the improved security would be reflected in major critical infrastructure systems such as nuclear power plant control; electric power generation and distribution; control of water distribution; building environmental, security, and safety controls; and manufacturing plant controls.
- The healthcare community, with NIST’s assistance, has used the CC for defining security requirements. Examples include: functional security requirements for Health Care Financing Administration’s Proposed Internet Security Policy; functional security requirements for the Department of Health and Human Services which maps the Health Insurance Portability and Accountability Act of

1996 Proposed Rule on "Digital Signature and Security Standards" into CC constructs; and a complete profile for patient "Point-of-Care Admission, Discharge and Transfer" in collaboration with Share Medical Systems (SMS).

As can be seen by these examples, the use of the CC for requirements specification is a first key step in improving the protection of our critical infrastructures—identification of sets of security requirements for IT products. This would have significant benefits even if security evaluations were not conducted. However, utilizing the CC as an evaluation tool against user-defined security requirements provides additional confidence that the products procured and deployed actually meet the desired security specifications.

The Road Ahead: Research and Resource Challenges

One of the criticisms often levied on NIAP is that evaluations take too long and cost too much. We hear this particularly from the small business community. Of course, one would expect to hear that of any evaluation process that is not free and instantaneous. But, in products involving great complexity and often millions of lines of code, such evaluations are time consuming. They also require rare expertise that is pricey in the marketplace. But we must ask ourselves whether improvements can be made? Indeed, given resolve, flexibility, resources, and research, I believe significant progress can be made.

Improving Current NIAP Testing

Here are some examples of what *could* be done:

- Develop NIAP guidance advising product developers how to reuse evaluation results from prior evaluations of the product.
- Develop NIAP guidance to maintain Common Criteria certificates for product maintenance changes (i.e., new versions) without the need to undergo a complete new evaluation.
- Develop an Assurance Maintenance module for the standard so only the changes to a previously evaluated product need be evaluated.
- Develop CC interpretations that clarify and simplify how parts of the CC are to be evaluated.
- Develop technology area-specific tests and test methods (e.g., smart cards, biometrics) that will provide more uniformity and comparability of evaluation results and result in more rapid evaluations for products.
- Using technology area-specific tests and test methods, establish accreditation criteria for labs that wish to specialize in evaluating products in a specific technology area (e.g., smart cards). Extend NIAP accreditation, on a voluntary basis, to those labs that wish to specialize in the technology area. This will result in cheaper, more rapid and more consistent evaluations for products in those technology areas
- Provide better training to lab evaluators and NIAP validators, with emphasis on which actions need to be performed and which do not.

- Provide an extensive/complete set of guidance documents for all stakeholders in the evaluation process (e.g., developers, evaluators, validators, commercial and government users).
- Provide clear guidance to stakeholders to choose only those assurance requirements that are meaningful for their intended use/environments.
- Perform a critical assessment of the current evaluation process to ensure that:
 - NIAP activities and levels of effort are consistent with those of other CC Recognition Arrangement partners
 - Evaluation activities are being performed efficiently
 - There are no unnecessary activities being performed
 - All activities that can be performed in parallel are in fact done that way.

We intend to seek out new partners, particularly in the homeland security community, to help support these activities in the near future.

Beyond NIAP

While these are key examples of what can be done to improve the current process, there is much more that should be done in order to address security assurance. Here are some examples:

- Conduct more research with the objective of developing new means to conduct security testing. The current techniques we have are either too expensive, involve too much human subjectivity, or both. The sooner the community pursues such research, the sooner we will benefit from their results.
- Develop comprehensive security requirements in both plain English and in the CC “language” that will be used to build more secure systems and networks. These security specifications must be developed with significant industry (users *and* vendors) and government involvement in key technology areas such as operating systems, firewalls, smart cards, biometrics devices, database systems, public key infrastructure components, network devices, virtual private networks, intrusion detection systems, and web browsers. These efforts can be adopted by voluntary industry consensus standards bodies as appropriate and can draw upon efforts underway in the NSA for national security systems.
- While it is important to understand and test security at the *product* level (the principal focus of NIAP), we need also to look outwards at the *system* and *enterprise architecture* level. For example, we need a means to rigorously understand the security implications that result when NIAP evaluated products are integrated together into a system. We also need to look inwards at IT building blocks such as protocols. Again, research will be a key to advancing our ability to make significant strides.
- We also need to look at other important security issues beyond just the (admittedly important) question of whether a product meets a security

specification. How do we gain assurance that the product does not do what is unintended? How can we gain assurance that no malicious code is buried deep inside software or hardware? How can we do such analysis as more and more development is taking place off-shore? Again, research is needed.

I would point out that the Cyber Security Research and Development Act of 2002 provides a means to support such research via academic and for-profit partnerships, in addition to intramural research at NIST.

Summary

The CC provides a means to develop security specifications and a common means to conduct security evaluations. NIST and NSA have created the NIAP, which uses accredited labs in the private sector to conduct such evaluation. However, more can be done to streamline this process through research and standards development.

Thank you for the opportunity to testify here today. I would be pleased to answer any questions you may have.

Mr. PUTNAM. Our second witness is Michael Fleming. Mr. Fleming currently leads the National Security Agency group responsible for development and customer implementation support of a broad set of IA solutions. Prior to this assignment, he held positions as the Deputy Chief of Network Security Group, Chief of Network Security Systems Engineering Office, Chief of Network Security Products Office and special technology transfer assignment with the NSA Deputy Director For Plans and Policy. Early in his NSA career Mr. Fleming served in a variety of technical and program management assignments in communications security and signals intelligence.

He is a recipient of the NSA Meritorious Civilian Service Award and twice received the Presidential Rank Award for Meritorious Service.

It is a pleasure to have you, and you're recognized for 5 minutes.

STATEMENT OF MICHAEL G. FLEMING, CHIEF, INFORMATION ASSURANCE SOLUTIONS, INFORMATION ASSURANCE DIRECTORATE, NATIONAL SECURITY AGENCY

Mr. FLEMING. Thank you for your interest in cybersecurity, information security, or information assurance. We have three words that describe this very important endeavor. I would like to provide a quick overview of the Common Criteria and the NIAP current status, some of the potential for even greater applications, and discuss some of the issues that you have already raised.

As you stated, it establishes a language which is very important, a syntax. The criteria is, in fact, a language, a dictionary for describing user needs and vendor claims. It also establishes the methodology to make those comparisons in terms of how well those claims meet needs.

I think it's important to note the criteria does not apply to all information technology products that make no information assurance claims. The criteria employs a distinct but related set of functional requirements which describe the mechanisms and the assurance requirements which Mr. Roback described in terms of gaining confidence that those mechanisms work correctly. There are seven assurance levels, one being the lowest and the least rigorous; seven being the highest and most rigorous.

In 1997, we entered a partnership with NIST called NIAP to promote, demand investment in security products, and establish the commercial security evaluation capability. To support the demand the Committee on National Security Systems in January issued NSTISSP 11, which stipulated the acquisition of commercial IA products, and IA-enabled products would be limited to those evaluated under formal schemes such as NIAP.

In terms of demand, we have defined, in fact, 21 protection profiles, and 31 more are in development. These profiles address key technology such as operating systems, firewalls, and intrusion detection systems and other things. And the demand trend there is encouraging.

As profiles are introduced for a technology, the number of evaluation claims is increasing. For example, all the operating systems in evaluation or that have been evaluated are compliant. All public key infrastructure are compliant. And about half of the firewall in-

trusion detection systems are either claiming compliance or compliant with protection profile.

As far as the second goal of NIAP, establishing the labs, 8 labs are accredited and have completed 38 evaluations with an additional 55 underway and more being negotiated continuously. In terms of expanding the use across a broader spectrum of environments than just the Department of Defense, the requirements for information assurance in the national security market are almost identical to those in other mission-critical government or commercial systems. Common Criteria can be leveraged to converge these markets. The larger market would result in greater return on investment for the vendors, and everyone in the buying sector would benefit from that leverage.

Regarding limitations, you address cost and timeliness. Evaluation of timeliness and cost actually is a function of a number of factors: Product complexity, assurance level aspired to, the vendor's preparedness to undergo an evaluation. And any problems found during evaluation typically want to be fixed before the evaluation is completed. All those can lead to some time and cost. But a vendor can capitalize on an initial evaluation investment by reusing parts for subsequent evaluations for subsequent releases.

While the criteria makes every attempt to identify and correct security vulnerabilities to ensure—there is no assurance that these products are bulletproof, especially at the lower assurance level. Vulnerabilities can be introduced in a number of ways, from poor design, inappropriate operation. Source code evaluation is not always required, particularly until you get to the higher assurance levels, which many vendors don't aspire to. And vulnerabilities in an IA-enabled product introduced by unevaluated nonsecurity functionality may go undetected. Mechanisms complementary to the Common Criteria are needed to increase our ability to find and eliminate malicious code in large software applications.

In conclusion, information systems require assurance that it was specified and designed properly, that it was independently evaluated against a prescribed set of security standards, that it will maintain proper operation during its lifetime even in the face of malicious attacks and human error. The Common Criteria in NIAP are working. The trends are up, and process improvements continue. A converged market for security products would benefit all potential IA buying sectors.

The Common Criteria and NIAP are not a panacea for all security issues and all information technology. We need complementary activities. Security needs to be baked into information systems starting with specification. It cannot just be evaluated in at the end nor sprinkled in after a system is fielded. And I think this is an important point in terms of improving the overall process. This is all about making sure that a security product is, in fact, secure and doing its job.

It has certainly been my pleasure to discuss the Common Criteria and share the work of the NIAP with the subcommittee, and I thank you for the opportunity.

Mr. PUTNAM. Thank you very much, Mr. Fleming.

[The prepared statement of Mr. Fleming follows:]



National Security Agency



Statement by

Michael G. Fleming

Chief, Information Assurance Solutions

Information Assurance Directorate

National Security Agency

Before The

House Committee on

Government Reform

Subcommittee on

Technology, Information Policy, Intergovernmental Relations

and Census

Hearing on

**“Exploring Common Criteria: Can it Ensure that the Federal
Government Gets Needed Security in Software”**

September 17, 2003

For Official Use Only
Until Released by the
U.S. House of Representatives
Committee on Government Reform

Thank you Chairman Putnam and the members of the Subcommittee. I am honored to have the opportunity to speak with your committee to discuss the Common Criteria and the National Information Assurance Partnership (NIAP).

I also would like to thank the Chairman and other members of the Subcommittee for their strong interest and attention to the vital area of cybersecurity. Your leadership is important for raising awareness of the serious security challenges we all face in our age of interconnected, inter-dependent digital networks.

My name is Michael Fleming and I am the Chief of the Information Assurance Solutions Group, Information Assurance Directorate, National Security Agency (NSA). My Group is responsible for developing information assurance solutions, support for the *International Common Criteria for Information Technology Evaluation* (known as the Common Criteria), and the NIAP.

I would like to note that the NSA's Information Assurance Directorate and its predecessor organizations have had technical and policymaking responsibility regarding the protection of national security telecommunications and information processing systems across the Executive Branch since 1953.

In regards to your theme for this hearing: "Exploring Common Criteria: Can it Ensure that the Federal Government Gets Needed Security in Software?" while in the security business it is hard to "ensure" absolutely, we believe the Common Criteria is a very important step in improving the "goodness" of an information assurance (IA) or information assurance enabled (IA-enabled) information technology (IT) product. I would like to provide you with an overview of the Common Criteria and the National Information Assurance Partnership (NIAP) and how it operates, highlight its benefits, and finally discuss the remaining issues associated with the activity. In Appendix A of my statement, you will find a synopsis of the lineage behind both the evolution of the criteria and the evolution of the evaluation programs for commercially produced IA or IA-enabled products to help understand the rationale behind the adoption of the *International*

Common Criteria for Information Technology Security Evaluation (subsequently referred to as the Common Criteria) and the establishment of the (NIAP).

The Common Criteria represents the outcome of an international effort (United Kingdom, France, Germany, Netherlands, Canada, and the United States) to develop criteria for the evaluation of information technology security by providing a standard language or syntax for describing the security requirements of an IA or IA-enabled product or system. Version 1.0 of the Common Criteria was published for comment in 1996, which was extensively reviewed and trialed by several nations. Based upon this review and lessons learned, the Common Criteria Version 2.0 was officially published in May 1998 and adopted by the International Organization for Standard (ISO) as an International Standard (ISO 15408) in August 1999.

For the purposes of this testimony and to put information technology products into perspective, I would like to categorize three types of information technology products; IA, IA-enabled, and other relevant IT products as shown in Figure 1: IA Relevant Technology Spectrum.




Figure 1: IA Relevant Technology Spectrum

An IA product's primary purpose is to provide security functionality (e.g., confidentiality, authentication, integrity, access control, or non-repudiation of data).

Examples of an IA product include Public Key certificate management, firewalls, intrusion detection devices, etc. An IA-enabled product is an information technology product whose primary role is not security, but which provides security functionality as an associated feature of its intended operating capabilities. Examples of an IA-enabled product include operating systems and database management systems with IA enabling functions (e.g., identification and authentication, passwords, audit, access controls, etc.), web browsers, e-mail, etc. Other relevant information technology products are those that provide no security functionality but do provide information processing services. Examples of other relevant IT products include switches, embedded software control modules, etc. This category is relevant because these products, while claiming no IA functionality, can be the source of vulnerabilities. An example would be the embedded timing module of a coolant system within a power plant.

One of the major benefits of the Common Criteria is that it establishes a common language for describing consumer security needs and IA or IA-enabled product vendor claims as well as the methodology for independently evaluating how well the claims meet the needs. While the Common Criteria is a very good specification and assessment tool for the security functionality within IA-enabled products, it should be noted that typically this functionality is only a subset of the total functionality of a product. As shown in Figure 2: Common Criteria Application to IA Relevant Technology Spectrum, the Common Criteria is applied to security functionality found in IA and IA-enabled products but is not applied to the functionality of other relevant information technology products since they make no IA claims. A Common Criteria evaluation typically analyzes the security functionality. Any vulnerability that is within an IA-enabled product that may be introduced by non-security functionality could go undetected (i.e., only the claimed IA functionality is typically evaluated).




Figure 2: Common Criteria Application to IA Relevant Technology Spectrum

The Common Criteria employs distinct but related categories of functional requirements and assurance requirements. Functional requirements describe security behavior mechanisms and assurance requirements describe the confidence gaining measures that the claimed security functionality is implemented correctly. For assurance requirements the Common Criteria defines seven (7) evaluated assurance levels (EALs). These EALs are denoted as EAL 1 through EAL 7 with EAL 1 being the lowest and least rigorous evaluation and EAL 7 being the highest and most rigorous evaluation. Further detail regarding the activities that are performed at each of the evaluation assurance levels is found in Appendix C.

International Mutual Recognition

Following the development of the Common Criteria, the authoring nations joined together to develop a Common Criteria Recognition Arrangement (CCRA). This recognition arrangement established the framework for each nation to mutually accept the validity of evaluations conducted by another nation for the first four evaluated assurance levels (EAL 1 through EAL 4) of the Common Criteria. Each member nation agreed that evaluations would be conducted using the Common Criteria and associated Common

Evaluation Methodology (the “how-to” companion document) to provide the member nations confidence that an evaluation would yield the same results regardless of which nation performed the evaluation. Mutual recognition of the product evaluation should not be construed as an endorsement, approval, or recommendation for use of the product by any member nation.

Establishment of the National Information Assurance Partnership

In September 1996, the NIST and the NSA entered into discussions on the creation of a joint testing center to focus on the evaluation of commercially produced IA or IA-enabled products against the emerging Common Criteria. These discussions were the genesis for the current National Information Assurance Partnership (NIAP). On August 22, 1997, the Director of NIST’s Information Technology Laboratory and the Deputy Director of NSA’s Information System’s Security Organization signed the formal Letter of Partnership. The partnership combines the extensive information technology security experience of both organizations to promote the development of technically sound security requirements for IA or IA-enabled products and systems and appropriate measures for evaluating those products and systems. The goal of the NIAP was to increase confidence in IA and IA-enabled products through independent, third party evaluation to help ensure the security of the information technology systems and networks. More specifically, NIAP sought to: 1) promote demand and investment in security products and 2) establish a commercial security product evaluation capability to compliment existing government evaluation and testing efforts. With the background set, lets now take a look at how well the NIAP is meeting its stated goals.

National Information Assurance Partnership Goal Achievement

The NIAP’s first goal was to promote demand and investment in IA and IA-enabled products. One of the major benefits of the Common Criteria is that it establishes a common language to describe consumer security needs and/or IA and IA-enabled

product vendor claims, as well as establishes the mechanism for independently evaluating how well the claims meet the needs.

In support of efforts to increase the use and availability of evaluated products the National Security Telecommunications Information Systems Security Committee (NSTISSC), which is now known as the Committee on National Security Systems (CNSS) issued NSTISSC Policy Number 11 (NSTISSP No. 11) in January 2000. The CNSS consists of representatives from 21 U.S. Government Departments and Agencies (listed in Appendix B).

NSTISSP No. 11 stipulates that information assurance (IA) shall be considered as a requirement for all systems used to enter, process, store, display, or transmit national security information. IA shall be achieved through the acquisition and appropriate implementation of evaluated and validated Government Off-the-Shelf (GOTS) or Commercial Off-the-Shelf (COTS) IA and IA-enabled Information Technology (IT) products. As of 1 July 2002, the acquisition of COTS IA and IA-enabled IT shall be limited to those products which have been evaluated and validated in accordance with the following:

- 1) The NSA/NIST National Information Assurance Partnership (NIAP) Evaluation and Validation Program,
- 2) The NIST Federal Information Processing Standard (FIPS) Cryptographic Module Validation Program, or
- 3) The International Common Criteria For Information Security Technology Evaluation Mutual Recognition Arrangement.

The acquisition of all GOTS IA and IA-enabled products shall be limited to those products which have been evaluated by the NSA, or in accordance with NSA-approved processes. The policy further stipulates that normally a complementary combination of IA and IA-enabled products are needed to provide a complete security solution to a given environment.

NSTISSP No. 11 did not stipulate specific security requirements from a functional or assurance point of view. The intent of NSTISSP No. 11 was to allow vendors to make claims about their products that could be validated and for consumers to decide if the validated requirements satisfied their needs. Paragraph 4 of NSTISSP No. 11 says; “it is important that COTS products acquired by U.S. Government Departments and Agencies be subject to a standardized evaluation process which will provide some assurances that these products perform as advertised.” By not stating any specific requirements other than evaluation, NSTISSP No. 11 gives vendors the flexibility make evaluatable claims about their product’s security functionality at a given assurance level using Common Criteria language that can be independently validated.

One of the major thrusts of the NIAP has been on using the Common Criteria as a way to state the security requirements that are needed by U.S. Government consumers in critical technology areas. The Common Criteria documents that state these security requirements are called Protection Profiles. Protection Profiles define an implementation-independent set of security requirements and objectives for a category of IA and IA-enabled products, which would meet the needs of a particular application environment. A Protection Profile has 6 sections that must be addressed so that it can be evaluated for conformance to the Common Criteria (see Appendix D).

Based on discussions with vendors and users (DoD and other Federal Government agencies), the NSA Information Assurance Directorate and the NIST have identified key IA and IA-enabled technologies and have undertaken efforts to define Protection Profiles for them. These key technologies include Operating Systems, Firewalls, Wireless, Web, Intrusion Detection Systems (IDS), Tokens, Databases, Virtual Private Networks (VPN), Biometrics, and Public Key Infrastructure (PKI). Currently, there are 21 finalized Protection Profiles of which eighteen (18) are U.S. Government and three (3) are from commercial organizations. Additionally, there are thirty-one (31) new U.S Government Protection Profiles under development.

DoD Directive 8500.1, "Information Assurance" and DoD Instruction 8500.2, "Information Assurance (IA) Implementation" characterize security application environments as needing low, medium and high security robustness. As such, the U.S. Government Protection Profiles state the security requirements necessary to protect information within the various security robustness environments.

The combination of these policy based demand incentives have been encouraging. As U.S. Government Protection Profiles are introduced for a particular technology sector, the number of evaluations claiming compliance with a Protection Profile has been increasing. For example 100% of all operating systems evaluations, 100% of all Public Key Infrastructure Certificate Issuing Management Components, 61.5% of all Firewalls, and 60% of all intrusion detection systems are claiming compliance or have met U.S. Government Protection Profiles.

The second goal of the partnership was to establish a commercially based evaluation and testing scheme to compliment existing government evaluation capabilities. The NIAP developed and established the policies and procedures for participation in the Common Criteria Evaluation and Validation Scheme and established the Common Criteria Evaluation and Validation Scheme Validation Body in 2000. This jointly staffed organization approves participation of commercial security testing laboratories in the scheme, provides technical guidance to those testing laboratories, validates the results of IA and IA-enabled product evaluations for conformance to the Common Criteria, and serves as an interface to other nations for the mutual recognition of such evaluations. Since its implementation the Common Criteria Evaluation and Validation Scheme has accredited, through the NIST sponsored National Voluntary Laboratory Accreditation Program (NVLAP), nine (9) commercial evaluation facilities, with eight (8) of these facilities still actively participating in the scheme to date. As of 31 August 2003, these facilities have completed thirty-eight (38) evaluations of IA and IA-enabled products. Additionally, there are currently fifty-five (55) IA and IA-enabled product evaluations currently on-going within the commercial evaluation facilities with these facilities

negotiating new evaluation contracts daily. These products, produced by large as well as small corporations, are from the spectrum of IA and IA-enabled products.

In order for the IA or IA-enabled product to be evaluated, the vendor of the product must develop a Common Criteria specification known as a “Security Target.” Unlike a Protection Profile, a Security Target is implementation specific. The Security Target contains all of the sections of a Protection Profile with an additional seventh section called the Target of Evaluation (TOE) Summary Specification. This section is where the vendor describes how their product satisfies the security requirements based on the environment, assumptions, policies, threats, and objectives.

Once a Security Target has been created, the IA or IA-enabled product vendor takes the Security Target to a NIAP approved or international mutually recognized Common Criteria Testing Laboratory (CCTL) for formal evaluation. Upon successful completion of the evaluation, a Common Criteria certificate is issued to the IA or IA-enabled product vendor and the Security Target and Validation Report are made available to the public (<http://niap.nist.gov/cc-scheme/ValidatedProducts.html>).

Aspects and Benefits of Criteria Based Evaluation

Along with the technology explosion comes a desire of the consumer to have confidence when they utilize their IA and IA-enabled products that their exposure to vulnerabilities are kept to a minimum. Even with a criteria based evaluation, no product can be deemed “Bullet-Proof.” Vulnerabilities can be introduced in a number of ways from product design and development, through poor implementation of their design, and through operation of the system. Vulnerabilities can be introduced into a product or system at the requirements definition phase if insufficient or ineffective requirements are incorporated into the product design. During the construction of the product, vulnerabilities can arise from incorrect design decisions or errors in design implementation. Once a product/system is installed, vulnerabilities can be introduced

due to inadequate controls or enforcement of these controls in the operational environment.

The question is how a criteria based evaluation can aid the consumer in mitigating most of the risks associated with using an IA and IA-enabled product. Being able to specify the needed security features (functionality) and the level of confidence (assurance) for IA and IA-enabled products is an important first step in building more secure systems. Using Protection Profiles provides manufacturers with a potential build to specification and a known potential market. Using an independent evaluation provides the consumer with a level of confidence that the vendor's claims are indeed valid. This confidence is gained through the various activities associated with an evaluation. The combination of activities and the rigor to which they would be applied will increase as the evaluation assurance level increases.

What are some of the Issues with the Common Criteria

The cost and timeliness of a Common Criteria evaluation varies depending on a number of factors: the complexity of the IA or IA-enabled product and the claims made in the Security Target; the Evaluated Assurance Level chosen (the higher the EAL the more likely the higher the costs); the vendor's preparedness to undergo an evaluation (vendors must provide specific documented evidence to support their claims); and problems found in conforming to the requirements must be fixed before the IA and IA-enabled product can complete evaluation. These costs are usually passed on to the consumer making evaluated IA and IA-enabled products more expensive than non-evaluated IA and IA-enabled products. However, the criteria and the NIAP evaluation program are structured such that a vendor can capitalize on their initial evaluation investment and re-utilize most if not all of their previous evaluation work to significantly reduce the cost and timeframe for subsequent evaluations of their next release at the same Evaluated Assurance Level or to migrate the evaluated product to a higher Evaluated Assurance Level.

While a criteria based evaluation makes every attempt to identify and correct security vulnerabilities and/or flaws within an IA and IA-enabled product from a security perspective given the size and complexity of most products and large number of lines of code, it cannot ensure that the product is “Bullet-Proof”, especially at the lower Evaluated Assurance Levels. The security functionality within an IA-enabled product is only a subset of all the functionality within the product. A Common Criteria evaluation will only analyze the security functionality at the selected Evaluated Assurance Level. Access to and evaluation of full source code is not required until the Evaluated Assurance Level 5, which is generally higher than most commercial vendors aspire to. Vulnerabilities within an IA-enabled product that are introduced by non-security functionality may go undetected. Historically, these vulnerabilities have been the most exploited. A significant cyber security challenge will be found in enhancing our ability to find and eliminate malicious code in large software applications. Beyond the matter of simply eliminating coding errors, this capability must find malicious software routines that are designed to morph and burrow into critical applications in an attempt to hide.

Applicability of Common Criteria Across Government and Beyond

The requirements for Information Protection and Information Assurance in our traditional national security market are almost identical to the IA requirements found in mission-critical government systems and the commercial critical information protection market. Many of these systems will be coming under the direct control or influence of the Department of Homeland Security. Legislation as recent as the Healthcare Information Protection and Privacy Act recognizes the need to protect and individual’s information.

We must accelerate the convergence of these markets and use the emerging Homeland Security policies to join these three communities into a single unified market for IA products. The unification on the demand side of the IA market will naturally result in greater interest on the supply side of the market to develop compliant systems. A larger market results in greater return on investment (ROI) for vendors, and everyone in the IA market benefits from the resulting reduced costs, increased functionality, and greater

assurance. A “converged market” for IA products market will also significantly increase the potential for interoperability among national security, mission-critical government, and critical infrastructure protection systems, to include similar systems operated by our international trading partners and military allies. The U.S. Government cannot afford to develop and deploy IA systems that do not interoperate or that require complex configuration or costly system management structures.

The Common Criteria and the NIAP evaluation scheme offer a mechanism for providing a standardized specification of these IA needs and an independent third party evaluation of a product’s conformance to these needs. Through the use of the NIAP evaluation program coupled with widely accepted Protection Profiles by the government and industry, a “converged market” could be created.

Conclusion

All information systems require the element of assurance. Assurance that the system was specified and designed properly. Assurance that it was independently evaluated against a prescribed set of explicit security standards. Assurance it will maintain proper operation during its lifetime, even in the face of malicious attacks and human error.

The Common Criteria and NIAP are working, the trends are up and process improvements continue.

A converged market for security products would benefit all buying sectors and the IA and IA-enabled product vendors.

The Common Criteria and NIAP are not a panacea for all security issues for all information technology. We need complementary activities. It has been my experience that security is most effective when it is “baked in” to information systems starting with

specification and continuing through design and development. Assurance cannot be “evaluated in” or sprinkled over a system after it is fielded.

It has been my pleasure to discuss the Common Criteria and to share the work of the NIAP with the sub-committee today and I thank you for the opportunity.

Appendix A

Evolution of Evaluation Criteria

A Defense Science Board Task Force report, “Security Controls for Computer Systems,” published in February 1970, made a number of policy and technical recommendations on actions to be taken to reduce the threat of compromise of classified information processed on remote-access computer systems. Department of Defense Directive 5200.28 and its accompanying manual DoD 5200.28-M, published in 1972 and 1973 respectively, responded to one of these recommendations by establishing uniform DoD policy, security requirements, administrative controls, and technical measures to protect classified information processed by DoD computer systems.

Concurrent with DoD efforts to address computer security issues, work was begun under the leadership of the National Bureau of Standards (NBS) (the predecessor to the National Institute of Standards and Technology (NIST)) to define problems and solutions for building, evaluating, and auditing secure computer systems. As an outgrowth of recommendations from this work, and in support of the DoD computer security initiative, the MITRE Corporation began work on defining computer security evaluation criteria that could be used to assess the degree of trust one could place in a computer system to protect classified data.

The National Bureau of Standards and MITRE evaluation material evolved into the *Department of Defense Trusted Computer Systems Evaluation Criteria* (also known as the Orange Book or DoD 5200.28-STD) which was released in 1983. It was later updated and re-released in December 1985 and served as the evaluation criteria for systems used within the federal government from 1985 until 2000.

In the late 1980’s Canada developed a similar criteria known as the *Canadian Trusted Computer Product Evaluation Criteria* (CTCPEC) and the European Community developed the *Information Technology Security Evaluation Criteria* (ITSEC). Each

established an accompanying evaluation program for commercial IA or IA-enabled product evaluation against the respective criteria.

In 1990, the NIST and the NSA launched an initiative to update the *DoD Trusted Computer Systems Evaluation Criteria* with a new jointly developed criteria for all of federal government known as the *Federal Criteria*. The Canadian and the European Community were also launching initiatives at this time to update their respective criteria. However in 1993, prior to the completion of the Federal Criteria, an international coalition of nations which included the United Kingdom, France, Germany, Netherlands, Canada, and the United States (NSA and NIST) reached agreement that a common security evaluation criteria should be developed rather than having a separate security evaluation criteria for each nation. The vendors of IA and IA-enabled products favored this approach because it would eliminate the need for three unique evaluations of the same product. This led to a pooling of international experts and resources directed towards the production of the *International Common Criteria for Information Technology Security Evaluation*. Version 1.0 of the Common Criteria was published for comment in 1996, which was extensively reviewed and trialed by several nations. Based upon this review and lessons learned, the Common Criteria Version 2.0 was officially published in May 1998 and adopted by the International Organization for Standard (ISO) as an International Standard (ISO 15408) in August 1999.

Evolution of Evaluation Programs

The National Computer Security Center, formerly named the DoD Computer Security Evaluation Center, was formed in January 1981 to staff and expand on the work started by the DoD computer security initiative.

The NSA through the National Computer Security Center implemented the Trusted Product Evaluation Program for the evaluation of commercially available computer systems against the *DoD Trusted Computer Systems Evaluation Criteria*. The Trusted Product Evaluation Program utilized government evaluators from the NSA and selected Federally Funded Research and Development Centers.

In December 1994, the NSA based on a NIST proposal and with their cooperation, took actions to implement a commercially based IA or IA-enabled product evaluation program. During this time of information technology explosion, IA and IA-enabled product explosion, and government downsizing, evaluation responsibilities shifted from a government funded and staffed evaluation program to a commercially-based, fee for service evaluation program. This action was essential if the U.S. was to maintain a viable program for the assessment of commercial-off-the-shelf (COTS) IA and IA-enabled products in a timely and cost effective manner. The decision for this fundamental shift was predicated upon the resource limitations of the government coupled with the lengthy timeframe for acceptance into and completion of an evaluation. After a two (2) year development and training effort, the NSA implemented the Trust Technology Assessment Program in January 1997, approving six commercial evaluation facilities to conduct evaluations against the *DoD Trusted Computer Systems Evaluation Criteria* with the IA or IA-enabled product vendor funding the cost of the commercial evaluation. The NSA continued to maintain oversight of each evaluation and issued the certificate of completion and compliance to the criteria.

Appendix B

Members of the Committee on National Security Systems (CNSS)

Department of State
 Department of Treasury
 Department of Defense,
 Department of Justice
 Department of Commerce
 Department of Transportation
 Department of Energy
 Office of Management and Budget
 Central Intelligence Agency
 Federal Bureau of Investigation
 Federal Emergency Management Agency
 General Services Administration
 US Army
 US Air Force
 US Navy
 US Marine Corp
 National Security Agency
 National Communication System
 Defense Intelligence Agency
 The Joint Chiefs of Staff
 Assistant to the President for National Security Affairs

Permanent observers represent the Defense Information Systems Agency (DISA),
 Department of Education, Federal Communications Commission (FCC), National
 Aeronautics and Space Administration (NASA), National Imagery and Mapping Agency
 (NIMA), National Institute of Standards and Technology (NIST), Nuclear Regulatory
 Commission (NRC), Chairman, Subcommittee on Information Systems Security (SISS),
 Security Policy Board Staff (SPB), National Reconnaissance Office (NRO), and the
 Critical Infrastructure Assurance Office (CIAO).

Appendix C**Evaluated Assurance Levels**

The activities used to gain assurance about an IA and IA-enabled product and the rigor to which they are applied increases as you move up the Evaluated Assurance Levels from 1 to 7. These activities include an analysis of the process and procedures used in the development of the product with a corresponding check to ensure that the process and procedures are/were being applied to the development of the product. An analysis of the requirements can be conducted to ensure they are sufficient and effective for the product's functionality and security purposes. These requirements can be further traced to the design representations to ensure they are reflected in the product design. The product can be analyzed to ensure that the actual product is reflective of the design representations thus insuring that all requirements have been implemented. Additionally, one can perform an analysis of the vendor's functional tests and test results to ensure that the product was adequately tested and yielded appropriate test results. The evaluation team could also perform their own independent functional testing as well as conduct penetration testing to see if they can break into the product or by-pass security mechanisms within the product. A flaw analysis of the product can be conducted in an attempt to insure that IA and IA-enabling feature flaws can be kept to a minimum. And lastly, an analysis of guidance documentation provided by the vendor can be conducted to insure that it adequately describes the IA attributes of the product and processes and procedures for appropriately utilizing them.

Various of these activities are applied to meet the following Common Criteria defined evaluated assurance levels.

- EAL 1 – Functionally tested
- EAL 2 – Structurally tested
- EAL 3 – Methodically tested and checked
- EAL 4 – Methodically designed, tested and reviewed
- EAL 5 – Semiformally designed and tested
- EAL 6 – Semiformally verified design and tested
- EAL 7 – Formally verified design and tested

Appendix D

Protection Profile Sections

A Protection Profile has 6 sections that must be addressed so that it can be evaluated for conformance to the Common Criteria. These sections are:

- 1) Security Environment – in this section the consumer describes the environment in which they would see this IA or IA-enabled product being used.
- 2) Secure Usage Assumptions – the consumer describes assumptions made about the IA or IA-enabled product in the areas of connectivity, physical locations, and personnel.
- 3) Organizational Security Policies - this section describes any organization security policies that the IA or IA-enabled product would be expected to enforce.
- 4) Threats to Security – the consumer identifies the threats that the IA or IA-enabled product is expected to address and the threats that the operating environment is expected to address.
- 5) Security Objectives - this section identifies the security objectives that should be achieved through the use of this IA or IA-enabled product.
- 6) Security Requirements – the consumer selects from Part 2 of the Common Criteria the functional requirements and from Part 3 of the Common Criteria the assurance requirements for which they would like to have an IA or IA-enabled product validated against.

Mr. PUTNAM. Our third witness for this first panel is Robert Gorrie. Mr. Gorrie is the National Security Agency integree serving as the Deputy Director of the Defensewide Information Assurance Program [DIAP], office, Office of the Assistant Secretary of Defense for Networks and Information Integration. Prior to his retirement from the Army after a 26-year career as a signal officer, he was Chief of the Information Assurance Division on the Joint Staff and Deputy Chief of NSA's Information Security Customer Support Office. Following his retirement, he is employed with Titan Systems Corp. as vice president of operations in its managed IT securities service group. He is a graduate of Gannon College and Penn State University in Pennsylvania as well as both the Naval and Air War Colleges.

Welcome to the subcommittee. You're recognized.

STATEMENT OF ROBERT G. GORRIE, DEPUTY DIRECTOR, DEFENSEWIDE INFORMATION ASSURANCE PROGRAM OFFICE, OFFICE OF THE ASSISTANT SECRETARY OF DEFENSE FOR NETWORKS AND INFORMATION INTEGRATION, AND DOD CHIEF INFORMATION OFFICER

Mr. GORRIE. Thank you, sir, and I am honored to be here and pleased to have the opportunity to speak with your committee about some of the efforts DOD has initiated with respect to the evaluation of information assurance and information assurance-enabled products.

As demonstrated in recent operations, U.S. forces have been extremely successful in the battlefield. They have been able to translate IT into combat power. However, as our dependence on IT increases, it creates new vulnerabilities as adversaries develop new ways of attacking and disrupting U.S. forces.

No one technology operation or person is capable of protecting the Department's vast networks. In October last year, the Department published its capstone information assurance policy. The policy establishes responsibilities and prescribes procedures for applying integrated, layered protection for DOD information systems and networks.

The DOD's IA strategies and policies are central to the committee's Common Criteria question. As I stated, no one single person, technology or operation can assure DOD's vast global networks. The Common Criteria, the NIAP evaluation program, the national and DOD policy addressing IA evaluations and the evaluated products themselves are part of an integrated DOD IA strategy.

Even with the solid defense-in-depth strategy in place, we must be confident in the security and trustworthiness of the products we use to implement that strategy. New vulnerabilities in the equipment we use are identified daily. Through the Department's IA Vulnerability Alert [IAVA], process, users are made aware of the vulnerabilities and associated fixes. The IAVA process serves us well, minimizing the effects of recent cyber incidents on DOD networks. The IAVA process has also highlighted the alarming rise in the number of vulnerabilities, the risks they represent and the cost of associated remediation. Although we continue to improve the efficiency and effectiveness of the IAVA process, unless we can take

proactive measures to reduce the number of vulnerabilities in our systems and networks, our ability to respond will begin to degrade.

Although no product will ever be totally secure, we can incorporate security into the design and, through testing, gain a reasonable sense of the risk we assume when we use them. However, that requires policy, enforcement and practice. The Committee on National Security Systems, their National Information Assurance Acquisition Policy directs the acquisition of all COTS IA and IA-enabled products to be used in national security systems be limited to those which have been evaluated. Our DOD policy goes further than that, requiring the evaluation of all IA and IA-enabled products.

While vendors are primarily driven by product cost, functionality and time to market, security has also become a significant consideration. Recently the largest vendors have pledged to make security a priority. The decisions of those vendors are based on thorough business cases analyses. None can afford the continued cost of the race against the “penetrate and patch” approach to deal with latent vulnerabilities in software packages. The economic cost of that approach is enormous and does not result in a higher level of security. Sound software engineering practices like those tested in a NIAP evaluation are an essential element in the elimination of vulnerabilities and critical to the reduction of postdeployment patching.

Still there remains the cost of evaluation and time of evaluation. Both are functions of the complexity of a product, the level of evaluation, the quality of a vendor’s product and the vendor’s preparation for evaluation. Product complexity in the evaluation level is directly proportional to the amount of testing required, and the amount of testing is directly proportional to the time and cost. A quality product may not require repeat testing. However, products that do get into a test, fix and test cycle incur additional cost not only for testing, but also for product modification.

Some vendors, especially small vendors, are concerned about the cost and time of evaluation regardless of the product’s complexity. During the development of DOD policy, we met with small businesses individually and in multivendor forums, and, based on their input, we developed policy that attempts to remedy some of their concerns.

The evaluation process does what it was designed to do. It provides standardized evaluation reports that help make—help us make informed risk management decisions with respect to the security of our networks and systems. Expectations of evaluated products should not exceed what the evaluations are designed to provide. The type of testing that uncovers vulnerability can be done by the NIAP laboratories and will be done if required. The depth of evaluation depends on how much time and how much money we are willing to pay, as well as how much risk we are willing to accept. Evaluations do not guarantee security. The security comes from sound systems engineering, the combination of technologies, operations and people.

The President’s recent National Strategy to Secure Cyberspace requires a comprehensive review of NIAP to examine its effectiveness and expansion potential. We are conducting that review in col-

laboration with the Department of Homeland Security. DOD is also investigating the issue of software assurance with respect to all software, not just IA and IA-enabled products, again working with the Department of Homeland Security.

The challenges we face are the same challenges found throughout government and industry, challenges we are addressing in our IA strategic plan. DOD is making progress managing the risks successfully across all of our national security and defense missions. That success is documented in our FISMA reports as well as our annual IA report to Congress. Most importantly, however, it's reflected in our ability to act as an enabler and not as an impediment in the conduct of networkcentric operations in several theaters across the globe.

I appreciate the opportunity to appear before your committee and look forward to your continuing support on this very critical issue. Thank you very much.

[The prepared statement of Mr. Gorrie follows:]

**Statement by
Robert G. Gorrie
Deputy Director
Defense-wide Information Assurance Program Office**

**Office of the Assistant Secretary of Defense for
Networks and Information Integration
and
DoD Chief Information Officer**

**Before The
Committee on Government Reform
Subcommittee on
Technology, Information Policy, Intergovernmental Relations and the Census
Hearing on
“Exploring Common Criteria: Can it Ensure that
the Federal Government Gets Needed Security in Software?”**

September 17, 2003

For Official Use Only
Until Release by the
Committee on Government Reform
U.S. House of Representatives

Thank you Mr. Chairman and members of the Subcommittee. I am honored to be here and pleased to have the opportunity to speak with your committee about actions the Department of Defense is taking to address threats to the security of its networks, systems and information. We continue to make significant progress in our quest to secure and defend our computer networks. My testimony will highlight some efforts we have initiated with respect to the evaluation of Information Assurance (IA) and IA-enabled products.

Secretary Rumsfeld, in one of his initial statements before the House Appropriations Defense Subcommittee, identified six key transformational goals for the Department. One of those transformational goals is to leverage Information Technology (IT) to create a seamless, interoperable, network-centric environment. As demonstrated in recent operations, U.S. Forces have unparalleled battlefield awareness; they can “see” the entire battlefield while the enemy cannot. They have translated IT into combat power beginning the transformation from Platform-Centric to Network-Centric Operations. And the transformation has just begun. A new era of warfare has emerged, one based on the concept that network connections provide greater power, agility, and speed. Multiple connections enable U.S. Forces to fight and mass combat effects virtually anywhere, anytime, and with a smaller “real” force. Through connections, smaller forces operating locally can leverage almost the full weight of global U.S. combat power. However, as our dependence on information networks increases, it creates new vulnerabilities, as adversaries develop new ways of attacking and disrupting U.S. Forces. In recognition of this relationship, the Secretary identified protection of U.S. information networks from attack as another of the transformational goals.

Secretary Rumsfeld describes transformation as an ongoing process, not an event – a journey that begins with a transformed “leading edge” force, which, in turn, leads the U.S. Armed Forces into the future. Mr. John Stenbit, Assistant Secretary of Defense for Networks and Information Integration and the DoD Chief Information Officer (CIO), is committed to support DoD transformation by providing the power of information to that leading edge. To bring power to the edge, he established the following goals: (1) develop a ubiquitous network environment, (2) populate the network with information of value, as

determined by the consumer, (3) ensure the network is highly available, secure and reliable. My role in bringing power to the edge is to support Mr. Stenbit's goals by guiding and overseeing the Department's Information Assurance (IA) Program; the strategy, policy and resources required to create a trusted, reliable network.

No one technology, operation, or person is capable of assuring or protecting the Department's vast networks and information. Everyone who uses, builds, operates, researches, develops and tests IT is responsible for assuring the Department's information and information infrastructure. A clear and coherent policy framework is required to ensure that individuals and organizations are aware of their responsibilities, and the Department's transformation to Network-Centric Operations is the framework we use to clearly define the "whys" and "hows" for such policy. For IA, net-centricity is a transformation of what we do, because the way we protect information and defend information systems and networks is fundamentally different in a globally interconnected world.

In October 2002, the Department published its capstone IA policy, DoD Directive 8500.1, "Information Assurance" followed in February the following year by amplifying policy in DoD Instruction 8500.2, "Information Assurance (IA) Implementation." The directive establishes basic policy and the instruction implements policy by further assigning responsibilities and prescribing procedures for applying integrated, layered protection of DoD information systems and networks.

The new policies establish a risk model to help information and system owners determine appropriate target levels of confidentiality, availability, and integrity. These target levels are expressed as IA Controls, which address security best practices for general threats and system exposures, federal and DoD policy requirements, and IA interoperability across the DoD Global Information Grid or GIG. The intent is to use these IA Controls as standard terms of reference for metrics and reporting. The Joint Staff has already taken a first step in that direction by cross-referencing them in the Joint Quarterly Readiness Review guidance, and we are working to make them the foundation of our Federal Information Security

Management Act (FISMA) reporting. DoD's Operational Test and Evaluation office will test the controls during the conduct of 'Red Team' assessments of newly deployed systems.

The DoD's IA strategies and policies are central to the Subcommittee's Common Criteria question. As I stated, no one single technology, operation, or person is capable of assuring DoD's vast global networks. The Common Criteria, the National Information Assurance Partnership (NIAP) evaluation program, the National and DoD policy addressing IA evaluations, and the evaluated products themselves are parts of an integrated DoD IA strategy. The technical strategy that underlies DoD Information Assurance is Defense-in-Depth, in which layers of defense are used to achieve a balanced overall Information Assurance posture. To take advantage of rapid advances in information technology the Department maximizes the use of COTS and balances this with layered security.

Even with a solid Defense-in-Depth strategy in place, a fundamental precept is our maintenance of confidence in the security and trustworthiness of the products we use to implement that strategy. New vulnerabilities in the equipment we use, both government and COTS, are identified daily. Through the Department's IA Vulnerability Alert (IAVA) process and attendant alerts, bulletins, and technical advisories, users are made aware of the vulnerabilities and associated fixes. The IAVA process serves us well, minimizing the disruption of DoD networks during recent cyber incidents that caused widespread disruption elsewhere. The IAVA process has also highlighted the alarming rise in the number of vulnerabilities, the risk they present, and the cost associated with their remediation. Although we continue to improve the efficiency and effectiveness of the IAVA process, unless we take proactive measures to reduce the number of vulnerabilities in our systems and networks, our ability to respond will begin to degrade.

The Chairman of the Joint Chiefs of Staff champions the concept of "born joint" as a way of expressing the need for built-in, seamless interoperability in new war fighting systems. Similarly, new IT products and systems must be 'born secure'; designed, tested, and validated against specific security requirements. The concept of 'born secure' combined with an aggressive vulnerability management program incorporating the IAVA process,

gives us the ability to proactively reduce our exposure to known vulnerabilities and maintain the capacity to respond to evolving vulnerabilities.

To help consumers select commercial off-the-shelf IT products that meet their security requirements and to help manufacturers of those products gain acceptance in the global marketplace, the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA) established a program under the NIAP to evaluate IT product conformance to international standards. The program, officially known as the NIAP Common Criteria Evaluation and Validation Scheme for IT Security, or Common Criteria Scheme in abbreviated form, is a partnership between the public and private sectors.

NIAP maintains a Validated Products List containing all IT products successfully completing evaluation and validation under the Common Criteria scheme. The validated products list also includes those products successfully completing similar processes under the schemes of authorized signatories to the Arrangement on the Mutual Recognition of Common Criteria Certificates in the field of IT Security. One of the challenges is to produce a full suite of U.S. security requirements, or protection profiles, required for industry to evaluate their products. The IA community is working hard to keep pace with the unique security requirements of constantly evolving and new IT by developing new protection profiles in collaboration with industry and academia.

Timeliness is a key performance parameter. The government must rapidly integrate secure cutting-edge products into its IT enterprise and industry must meet time-to-market requirements. We cannot still be evaluating Version 4.0 of a product when Version 6.0 is on the market. In the aftermath of the events of September 11, NIST and NSA accelerated the protection profile development process and recently announced a new collaborative effort to produce comprehensive security requirements and security specifications for key technologies that will be used to build more secure systems for our Federal Agencies. These security requirements and security specifications will be developed with significant industry involvement. Protection profiles in key technology areas such as operating systems, firewalls, smart cards, biometrics devices, database systems, public key infrastructure

components, network devices, virtual private networks, intrusion detection systems, and web browsers will be the primary focus of this high priority project. With defined product security requirements and specifications, a defined and efficient product evaluation process and most important, a strong partnership with industry we will be able to populate the Validated Products List with up to date and secure IA and IA-enabled products.

Although no product will ever be totally secure, we can incorporate security into their design and through comprehensive security test and evaluation gain a reasonable sense of the risk we assume when we use them. However, for that concept to become a reality, it must be codified in policy and enforced in practice. In January 2000, the Committee on National Security Systems (CNSS), formerly the National Security Telecommunications and Information Systems Security Committee, issued its National Information Assurance Acquisition Policy. That policy directs, "by 1 July 2002, the acquisition of all COTS IA and IA-enabled IT products shall be limited only to those which have been evaluated and validated in accordance with criteria, schemes, or programs of the Common Criteria, the National Information Assurance Partnership (NIAP) evaluation and validation program, and the Federal Information Processing Standards (FIPS) validation program."

DoD policy goes further than the National policy, requiring the evaluation of all IA and IA-enabled products, not just those used in National Security Systems. Department acquisition policy includes references to the mandates of CNSS and DoD IA policy to insure IA is a key element of all acquisitions. The combination of the CNSS and DoD policies, the Common Criteria IA validation scheme, and the development of Protection Profiles in key IT areas is the foundation for 'born secure' IT.

Internal to the Department, Services and Agencies have published supporting service/agency specific policy for the evaluation of IA and IA-enabled products. We have an aggressive NIAP awareness campaign within the department. We also have enacted controls to monitor and enforce compliance with policy. The first conversations between a vendor and user often center on the requirement and timeline for NIAP evaluation.

While vendors' drivers are primarily product cost, functionality and time-to-market, security has become as significant consideration. Recently, the nation's largest vendors have pledged to make security a priority. For example, on Jan 15, 2002, Bill Gates released an email stating Microsoft's highest priority. *"Trustworthy Computing is the highest priority for all the work we are doing. We must lead the industry to a whole new level of Trustworthiness in computing."* Microsoft's decision and the decision of many other vendors to focus on security are based on thorough business case analyses. None can afford the continued cost of the race against the "penetrate and patch" approach to deal with latent vulnerabilities in software packages. Simply, the economic cost of this "penetrate and patch" approach is enormous and does not result in a higher level of security. Sound software engineering practices, like those tested in a NIAP evaluation, are an essential element in the elimination of vulnerabilities and critical to the reduction of post deployment patching.

Still, there remains the cost of evaluation and the time of evaluation. Both are functions of the complexity of a product, the level of evaluation, and the quality of a vendor's product and preparation for evaluation. The amount of testing required in evaluation is directly proportional to product complexity and evaluation level. The amount of testing relates directly to time and cost. A quality product will not require much repeat testing. Products that get into a test, fail, fix, and test cycle incur additional costs not only for testing but also for product modification.

Some vendors, especially small vendors, are concerned about the cost and time of evaluation regardless of product complexity and evaluation level. During the development of DoD policy, we met with small businesses, individually and in multi-vendor forms. Based on their input, we developed policy that attempts to remedy some of their concerns, specifically the concern over the investment in evaluation without knowing if there would be a return on that investment. e.g., DoD policy states, "...products must be satisfactorily evaluated and validated either prior to purchase or as a condition of purchase; i.e., vendors will warrant, in their responses to a solicitation and as a condition of the contract, that the vendor's products will be satisfactorily validated within a period of time specified in the solicitation and the

contract.” Vendors can now enter competition and if selected realize a return on their evaluation investment. Other modifications were also made to policy based on consultation with industry.

Questions have been raised about the efficacy of the end-to-end evaluation process itself and the extensibility of the process to the entire Federal government and civil community beyond National Security System users. The evaluation process does what it was designed to do. It provides standardized evaluation reports that help us make informed risk management decisions with respect to the security of our networks and systems. Expectations of evaluated products should not exceed what the evaluations are designed to provide. If a protection profile at a particular evaluation level does not call for the evaluation of some security functionality, it will not be evaluated. The type of testing that uncovers vulnerabilities like the buffer overflows exploited by some of the recent worms can be done by the NIAP laboratories and will be done if required. The depth of evaluation depends on how much time and money we are willing to pay as well as how much risk we are willing to accept. Evaluations do not guarantee security. The security comes from sound system security engineering, the combination of technologies, operations and people.

The President’s recent “National Strategy to Secure Cyberspace” requires a comprehensive review of NIAP to examine its effectiveness and expansion potential. We are conducting that review in collaboration with the Department of Homeland Security (DHS) to support the President’s strategy as well as the need for the evaluation process to keep pace with technology and DoD’s overall transformation efforts. DoD is also investigating the issue of Software Assurance with respect to all software, not just IA and IA-enabled products, again working with DHS. Our review of NIAP will help us improve the process and incorporate changes that will give us more confidence in the security of our IA and IA-enabled products.

The challenges we face are the same challenges found throughout government and industry – challenges we are addressing in our IA Strategic Plan. Does DoD have unique challenges – yes, but they are not insurmountable. Size, global presence, dynamic technical and operational requirements all contribute to the complexity of the Department’s environment.

But, DoD is making progress, managing the risk successfully across all of our National Security and Defense missions. That success is documented in our FISMA reports as well as in our Annual IA report to Congress. Most importantly, however, it is reflected in our ability to act as an enabler, not an impediment, in the conduct of Network-Centric Operations in several theaters across the globe.

We have come to realize that we will never be able to achieve absolute protection of our information, systems and networks. However, we also realize that we can effectively mitigate the effects of challenges to the security of our information, systems and networks. We have created a robust Computer Network Defense capability within the Department, a capability that continues to evolve and transform itself in pace with the evolving and transforming threat.

IA is a journey, not a destination. That may be a trite phrase but it accurately depicts the IA environment in DoD. Most systems are legacy systems as soon as they go online. The demand for greater bandwidth, functionality, connectivity and other features is constantly expanding. The IA challenge within the Department is to insure it is met securely. IA must be 'baked in' and not 'spread on' as an afterthought. DoD and the DIAP are stepping up to that challenge. DoD's IA community is intimately involved not only in the development of protective technologies for space-based laser, advanced fiber optic, and wireless transport networks but also in the development of end-to end IA architectures and technologies. From the labeling of information and people for controlled access to the security of enterprise computing environments, we are working now to ensure IA is 'baked in' and products are 'born secure' from both the protect and defense perspectives.

I appreciate the opportunity to appear before the Subcommittee and look forward to your continuing support on this very critical issue. Thank you.

Mr. PUTNAM. Thank you very much. And we are delighted to have been joined by the ranking member of the subcommittee, the gentleman from Missouri Mr. Clay, and the distinguished gentlelady from California Ms. Watson. And at this time I will recognize the ranking member for his opening statement.

Mr. CLAY. Thank you, Mr. Chairman, and especially for calling this hearing.

I'd like to reiterate two points that I made at last week's hearing. First, the government should use its power in the computer software marketplace to acquire safer software. Second, software vendors should be more aware of the security configuration of the software they produce. Let me briefly elaborate on these two points.

The Federal Government spends billions each year on computer hardware and software. Those purchases have a strong influence on what gets produced and sold to the public. The Federal Government can use its market power to change the quality of software produced by only buying software that meets security standards. The result will be an increase in the security of all software and better protection for the public. This is a simple formula. The government doesn't have to regulate software manufacturers, it only has to use its position in the marketplace.

Mark Forman, the former Federal CIO and regular witness before this subcommittee, incorporated an idea similar to this when he developed the Smart-Buy program. Mr. Forman realized that Federal agencies were buying the same software over and over again. Each agency was paying a different price for the same software, and the Federal Government was getting little or no leverage out of its position in the marketplace. No business would operate like that.

I believe we should build on Mr. Forman's idea to buy not cheaper software, but better software. I hope the new CIO, Karen Evans, will work with the subcommittee to incorporate this concept into the Smart-Buy program. We don't have to wait for computer companies to develop new security procedures. There are some steps that can be taken very quickly to improve computer security. We saw this earlier this year when Microsoft began shipping software that was configured differently.

The story Microsoft tells is that the company realized that it was shipping software with all the gates opened. A good computer manager systematically went through the software, closing gate after gate. Those with less training left the gates open, and the hackers walked in.

Shipping software with secure configurations should be a first priority of all computer companies.

I look forward to the testimony today of these witnesses, and I hope that our witnesses will consider my suggestions and provide the committee with their comments on it.

Thank you, Mr. Chairman, for yielding.

Mr. PUTNAM. Thank you very much.

[The prepared statement of Hon. Wm. Lacy Clay follows:]

**STATEMENT OF THE HONORABLE WM. LACY CLAY
AT THE HEARING ON
COMPUTER SECURITY**

SEPTEMBER 17, 2003

Thank you Mr. Chairman for calling this hearing. I would like to reiterate two points that I made at last week's hearing. First, the government should use its power in the computer software market place to acquire safer software. Second, software vendors should be more aware of the security configuration of the software they produce. Let me briefly elaborate on these two points.

The federal government spends billions each year on computer hardware and software. Those purchases have a strong influence on what gets produced and sold to the public. The federal government can use its market power to change the quality of software produced, by only buying software that meets security standards. The result will be an increase in the security of all software, and better protection for the public.

This is a simple formulation. The government doesn't have to regulate software manufactures. It only has to use its position in the market place.

Mark Forman, the former federal CIO and regular witness before this Subcommittee, incorporated an idea similar to this when he developed the Smart-Buy program. Mr. Forman realized that federal agencies were buying the same software over and over again. Each agency was paying a different price

for the same software, and the federal government was getting little or no leverage out of its position in the market place. No business would operate like that.

I believe we should build on Mr. Forman's idea to buy, not cheaper software, but better software. I hope the new CIO, Karen Evans, will work with the Subcommittee to incorporate this concept into the Smart Buy program.

We don't have to wait for computer companies to develop new security procedures. There are some steps that can be taken very quickly to improve computer security. We saw this earlier this year when Microsoft began shipping software that was configured differently.

The story Microsoft tells is that the company realized that it was shipping software with all the gates open. Good computer managers systematically went through the software, closing gate after gate. Those with less training left the gates open, and the hackers walked in.

Shipping software with secure configurations should be a first priority for all computer companies.

I look forward to the testimony today, and I hope that our witnesses will consider my suggestions and provide the committee with their comments on them.

Thank you Mr. Chairman.

Mr. PUTNAM. At this time we will recognize the gentlelady from California Ms. Watson for her remarks.

Ms. WATSON. Thank you so much, Mr. Chairman. I do appreciate this opportunity.

Over the last decade or so, the Internet has become a force in our society that it is difficult to identify critical networks in our Nation that are not connected to the Internet. Electricity, traffic, water, freight, all these systems rely on the Internet for their function. This reliance on the Internet has yielded tremendous gains in efficiency, yet we are constantly reminded of the vulnerabilities inherent in such reliance on the Internet.

Most recently, the Blaster and the SoBig viruses posed major challenges to the integrity of America's infrastructure. Thankfully, none of the cyber attacks known to us have resulted in cataclysmic damage to the United States, or to our people, or to our infrastructure, at least not yet. We have had many close calls. And in the wake of September 11, many analysts familiar with global terrorism blame America's leaders for missing the signs that we were vulnerable to conventional terrorism. If we in Congress do not wake up to the clear warning signs of our vulnerability, we would be committing just as grave a mistake.

In my experience in Micronesia, in my embassy, is that we were getting warnings by cable from the State Department on a daily basis of a virus that ran through our most sensitive computers in the embassy. That's a very scary notion when you depend on the Internet 24/7 to communicate.

And so this hearing is very, very valuable to the basic security of our country, and I really would like to be here to hear every bit of the comments that are being made by the panel with such expertise. But we have a hearing on terrorism, and I do hope our enemies around the globe do not—are not able to master the Internet to the extent that they know more than we do and they can get our country's secrets.

So thank you so much, panelists, in bringing your expertise to us. Thank you, Mr. Chairman. And I'm going to run down to that classified hearing.

Mr. PUTNAM. You just throw us off like a bad habit.

Ms. WATSON. I want to find out what those real secrets are.

Mr. PUTNAM. We will begin with the questions for the first panel.

Mr. Gorrie, could you explain why DOD decided to adopt the Common Criteria requirement for all DOD procurement. What led to that decision?

Mr. GORRIE. The original NSTISSP 11 requirement was for—only for national security systems, and if you look at the term “national security systems,” that's a legislative construct that was brought into being in I believe it was the Nunn-Warner amendment to the Brooks Act. The Brooks Act established that all ADPE, automatic data processing equipment, would be bought through GSA. The Department of Defense found that GSA wasn't really responsive to that. This was in 1986. And the Warner amendment, as it was known, changed that to reflect the term “national security systems,” so it said that—and I have it here somewhere, but—and I'll get it for you later, but it says all national security systems—and it went on to list what a national security system was: anything

that handled classified information, did intelligence, did cryptologic work, did weapons systems, were national security systems, with the exception of what they term support systems, which were personnel systems, logistics systems and things of that nature.

If you went out and talked to any commander around the globe and asked them if their personnel system or their logistic system was a critical part of their warfighting capability, they would undoubtedly all say yes, and so that is why—the reason we in DOD said you need not only security in your national security systems, but in all other systems that we use, because they all touch one another, and the potential for a security flaw in one could spill over to other ones.

Mr. PUTNAM. How would you evaluate your experience thus far in terms of the weaknesses of it, the strengths of it, lessons we can derive as we contemplate its usage beyond DOD?

Mr. GORRIE. First, the weaknesses of it. A lot of it has to do with our interaction with our vendors. Some of the vendors are not—and even some of the people, the users in DOD who have to follow the rule, are not familiar exactly with what the rule entails. Some of the criticisms from small businesses that they can't realize a return on investment are borne out of ignorance of what the policy says, because the policy does provide them the opportunity, when making a contract with the government to sell their particular product, that the only thing that they need to do is to stipulate in the contract that they will have the product evaluated, not that they have to evaluate the product prior to establishing the contract, which gives them the opportunity, if selected, to realize a return on that investment because they can include that cost in the cost to the government.

The number of systems which are being evaluated, although adequate right now, needs to be much, much higher, and the types of systems that are being evaluated need to be expanded. Those are our problems.

Benefits, as I said in my testimony, the ability to know what a product will do is one of the biggest benefits we can have. You can get the glossy brochure from a vendor that says this is the best thing since sliced bread, but until you put it to the test, you don't know what that product will do. And an independent evaluation such as that provided by the NIAP is invaluable not only because you know what it will do, but when you certify and accredit a particular system to be able to be connected to our networks, you have to make a risk decision whether or not that system is safe enough. If you know exactly what those products are doing, then you can craft other things around that particular product to circumvent any shortcomings it may have, things like an operational procedure or some kind of policy control or other things. So in that particular sense, the reports that we get out of NIAP are invaluable in order to make our systems safe.

Mr. PUTNAM. Thank you.

Mr. Fleming, when you developed NSTISSP 11, the requirement that national security systems purchase software certified after the Common Criteria, what consideration was given to its impact on small business?

Mr. FLEMING. NSTISSP 11, first of all, comes out of the national security community, which comprises of some 21 or 22 Federal departments and agencies and sort of the national security slice across those agencies, where in DOD it might go deeper than it would in some of the other agencies. It also requires an evaluation of all information assurance products.

The NIAP process is only one of the schemers. NSA does evaluations for high-grade cryptography. So the NSTISSP 11 applies to a broader thing than just NIAP.

As far as small businesses are concerned, the cost of evaluation, as I mentioned in the testimony, varies considerably depending on the assurance level. And when NSTISSP 11 was originally issued, it did not specify that all products had to be evaluated in the beginning. It put a date in there of July 2002. It came out in 1999. There was a period in there where it was something to be considered. And the idea there was to allow companies to get used to both the process and the profiles that were coming out. So the mandate did not start until, in fact, almost 2 years later in 2002. So the idea was to allow companies to grow toward what this was.

The second thing was it didn't specify any particular evaluation level. The beginning thinking was any evaluation level is better than none. And so the cost is, in fact, considerably lower at the lower levels than it would be at the higher levels because of the demand for generating evidence. So the idea was to ramp this process up to allow companies to grow with it and, over time, ever increase the assurance level in these products.

So that's how we wanted to consider, in fact, all vendors, but in particular the small companies.

Mr. PUTNAM. In the beginning of your answer, you mentioned that this was to cut across national security systems. Does the Justice Department and Homeland Security and State also utilize government criteria?

Mr. FLEMING. Yes. NSTISSP 11 includes all those agencies and the opportunity for them—obviously NSTISSP 11 applies at that level the opportunity to use the Common Criteria, and the NIAP process is there for any buyer. But, yes, NSTISSP 11 covers those kinds of agencies for their national security systems.

Mr. PUTNAM. Mr. Roback, as all of us are aware, and we have held hearings related to this, the Blaster worm exploited a flaw in Microsoft's operating systems to infect thousands of computers. Since that system was certified, why wasn't the flaw found? What is the weakness in the evaluation that does not get at code flaws?

Mr. ROBACK. I think you have to look at the range of possibilities that the NIAP testing program offers. At the low end, where you are looking at things like documentation of how the product was developed, you are not getting into the very detailed code review that you get at the very, very high levels of assurance. So it sort of depends on what level you want to pick for your evaluation, which is the flexibility of it. A vendor can bring in product and target any one of the seven levels or create their own. So unless they target something at the very high level, which, by the way, costs a lot more and takes a lot longer, you are not going to get that level of review. And even if you do, it's subject to human subjectivity in the review.

So you may not get—because we don't have very specific standards for this, and you probably couldn't at that level for millions of lines of code—standards you can do very quick, very exact testing. So there's some art in here, too.

Mr. PUTNAM. Thank you very much.

Mr. Clay.

Mr. CLAY. To all of the witnesses, I would just like to hear from you or hear your comments on the proposal to add secure configurations as another dimension of a Common Criteria. Is this feasible, and how long would it take? We'll start with you, Mr. Roback.

Mr. ROBACK. Actually under the Cybersecurity R&D Act that was passed by the Congress late last fall, they assigned to NIST the task of developing security configurations for specific IT products. And so we are holding a workshop later in September trying to invite the vendors in, and other Federal agencies and NSA and others have already developed some of these checklists for some very specific products. So some of these do exist.

Actually I think it would be a very good thing, because if you look at a spectrum, first you want to have very strong standards. Then you want to have some testing program that tells you whether the standard was correctly implemented. And third, you want to have those configurations so that when a system gets one of those products, they know where to set the settings, because even if they are shipped from the vendor with security turned on, which is not always the case, but sometimes it is, it is not necessarily always right for the environment that it's being put into.

Configuration guidance is a very good thing. It's also important to remember there's a range of potential environments; that is, the security you would have for a home user might be very different from the security at NSA or the security of a large, centrally managed enterprise. So you have to keep that in mind, too, there's a range there, because there's a range of risks in the type of information that's being exposed.

So it does get complicated, but I think checklists of that sort are very useful.

Mr. FLEMING. I would agree with everything Mr. Roback said. This is a life cycle. Security is a life cycle endeavor. It just doesn't stop when the product is certified and goes out into the field security. Every day you've got to watch these products, particularly security products that sit sometimes in the way of system performance, and it is so often tempting to tweak that firewall a little bit to allow the bandwidth to get greater, but you may have, in fact, left open the door you don't want to open.

So I would add to Mr. Roback's points the human dimensions of this. It boils down to how well trained is that system administrator or that security administrator; how well do they understand the multitude of configurations that these products can, in fact, take, and which ones are the good ones and the bad ones. So there's a dimension in this of awareness and training of individuals along with the ideas that Mr. Roback put forth in terms of having configuration guides. And we have been a very, very strong partner in the generation of these configuration guides for major IT systems, but there are many other technologies that need a similar kind of

guide for a well-meaning, but sometimes difficult job called system administration, security administration.

Mr. CLAY. Mr. Gorrie, anything to add?

Mr. GORRIE. Well, I could attest that it does work. In DOD we have been using secure technical implementation guides [STIGs], for our products for our operating systems and other things for a long time. We have a process that is known as gold disk, where we will go out and put particular security settings on operating systems.

STIGs, security technical implementation guidance. We have them in DOD. They work, and it depends on, again, as Mr. Roback stated, what environment you want to use them in. If you are using them for an inventory system in a gym, no sense in tightening it all down because you want to be more open and share those sorts of things. If you have a critical system that you need protected, it needs to be ratcheted down. And all the people who participate in that network have to have it ratcheted down to the same degree.

Mr. CLAY. In your opinion, would it be possible to certify software configurations separate from the Common Criteria evaluation?

Mr. GORRIE. I don't know. I would have to defer on that.

Mr. CLAY. Anybody?

Mr. ROBACK. Well, I am not sure if certification is the precise word, but I think that there are indeed—you can separate the two. Whether a product has been tested to know whether the security features work correctly is a separate question from where to turn on and turn off the security features.

However, if you haven't gone through certification, the testing process, you are not going to have a great deal of assurance that even if you turn something on, the security is working. And the example I like to give is you go to a Web site, and you get the little lock in the corner on your browser. Well, why do you have any confidence whatsoever that it is doing anything other than showing you a little picture of a lock? If it hasn't gone through testing, you really don't know, other than it makes a nice little picture in the corner.

So that is why testing is so important in addition to turning on the security.

Mr. CLAY. For all of the witnesses, again, I would like each of you to comment on my proposal that the Federal Government use its market power to improve the level of security for our purchasers. Do you believe this is feasible?

Mr. FLEMING. I will start. During the testimony I used a phrase called converged market, and I think it is along the lines of your reference back to the Smart Buy Program that Mr. Foreman put in place.

The idea of a converged market would be find that level of security goodness, that assurance level and that set of security mechanisms that a large buying sector could agree to, the DOD, the national security community, the other Federal agencies, the critical infrastructure marketplace that Ms. Watson referred to, such that a vendor would see a return on investment good enough for them to shoot for that level.

And so this idea of getting a common level that all would buy into would, I think, be a good incentive for vendors. Make it appropriate so it is not a bridge too far, and then standardize on that level and let vendors shoot for that level so you can get this economy of scale.

Mr. CLAY. Thank you.

Mr. Gorrie, one question from our other Member who had to leave. She says: It is good to hear that you understand the business costs of the reactive plug and patch approach, but how widespread do you feel this view has been accepted throughout the technology industry? What can we do to spread this message and change the approach?

Mr. GORRIE. I think if you will ask the panel members that follow us, those are the words that were given to me by them. I mean, they were the ones who told me those things. I didn't make that up.

How can we spread it? I think it will spread itself. Vendors whose products are well developed and have fewer problems as far as having to go out and patch them and things of that nature will be bought more. People will see the benefit of buying them, not being able to be hacked, not having to go in and reengineer their systems every time a patch comes out.

Those who have products which are constantly being patched will find their position in the marketplace becoming lower. It is a self-regulating system, and it will become more so in the future as more and more patches have to be made to accommodate shortcomings in software.

Mr. CLAY. OK. I thank the panel.

Thank you, Mr. Chairman.

Mr. PUTNAM. Thank you, Mr. Clay.

Mr. Fleming, under Common Criteria evaluation, the product is tested by itself. Obviously it will be used in conjunction with a variety of other products. Is that taken into consideration at all? And how is that issue resolved in terms of the impacts or the problems that can occur with the connectivity?

Mr. FLEMING. Good. First of all, the product is typically tested in an environment. In order to make the test meaningful, there is an assumed environment. However, that may not represent all possible environments for the application in the real world.

So there is—and these are my terms. There is this “little c” certification, which is certifying that the product is doing what its claim is. Then there is the application of that product, along with other products, into a larger system. So there needs to be another certification, which is at what I will call the “big C,” at the system level. There are processes in place in the DOD and across the Federal communities that go by somewhat different names, and they are like kind of complicated names like DITSCAP, Defense Information Security Accreditation Program, but that is a system-level look. So I see the Common Criteria and any other evaluation program at the product level generating evidence about the performance of the various components. Then there needs to be a separate look at the much larger level, for what the total system security certification is.

Now, I will state that the calculus for that is somewhat difficult, because it is not just saying product A has this level of goodness, product B has this level of goodness. You just don't add A and B and get C. In fact, it is a much more complex relationship when you start bringing many products together. But, nevertheless, there are processes in place in the DOD and beyond the DOD to bring this larger certification into play toward the ultimate accreditation of that system to operate in a real environment.

Mr. ROBACK. If I can just add to that, that question of when you understand the property of one component and then the property of another, and you put them together, that is what the researchers call the composability issue, trying to understand in a rigorous way what you have when you put them together and add them up.

If I could just add to Mr. Gorrie's comment earlier about the software quality and patching, I think one of the problems we face is the whole developmental cycle in the industry of software products. And you have to really look at how software products are developed in a rigorous sense of specifications and so forth.

If you really want to improve the overall security and get away from this problem of continually chasing our own tail and trying to patch, this is a Web site where we put up vulnerabilities in commercial products so people can learn about them and go find fixes for them. Right now we have over 6,000 vulnerabilities up there.

And, you know, the tolerance of the marketplace for these products that come out with flaws is just astounding. We really need to look at the overall quality issue of the products; not just the security, but the overall quality. Do they do what they are intended to do? It is a challenge.

Mr. PUTNAM. Mr. Gorrie, several testimonies mentioned a waiver process for the Common Criteria. Under what circumstances would a waiver be granted?

Mr. GORRIE. With the way that we have constructed policy within the DOD, I would find it very rare that you would have a request for a waiver. The only one that I could think of would be in a situation if we were going to war, we needed a particular product because its security features were just so obviously great that we could not not afford to have it in our system. But even then, the way that we have policy built, it says that you need not to contract to purchase that piece of equipment, you need not have it evaluated, only have it in the contract that you will have it evaluated as a condition of purchase.

However, the vendor always has the option to say, you may need this, this may be the best thing since sliced bread, but we are not going to have it evaluated. If we needed that product that bad, then the user of that product, or the person who wanted to put that product into their system, would have to petition for a waiver, and then we would either have it evaluated internally within DOD through some process, or just use it because it was so important to use. But in any regular process I—because of the way policy is written, I do not see the need for waivers.

Mr. PUTNAM. My final question for this panel will be this, and we will begin with Mr. Roback: Should the Common Criteria certification be extended to cover the entire Federal Government?

Mr. ROBACK. That is a good question, one we are often asked. Let me just start by mentioning that it is policy for the nonnational security side of government that cryptographic products have to go through testing, and there is no waivers allowed for that under FISMA.

I don't think the question is necessarily should we adopt that policy, yes or no. There is actually quite a range of options between doing nothing and adopting that policy, and even things beyond that policy. So you might ask yourself: Well, maybe it doesn't make sense to say something can be certified against any specification that is brought forward, but maybe what we need to do is look at things like once we have good specifications for specific technology, that if an agency is buying that technology, they should be buying something that has been evaluated against those specs.

So not just that you can bring in—I think someone in their testimony talked about a product that paints the screen blue, and it can go through and get a certification. Well, I don't know if those products are going to do us any good. So I think there is some range of options we have here, and we really need to look at those. Rather than just say, that is the policy for national security; we should simply adopt it.

I think we need to learn more from the experience as well. Is it really pushing the vendors toward more security or not?

Mr. PUTNAM. Mr. Fleming.

Mr. FLEMING. We are putting our trust in networks, in things called security products. They have become sort of a foundation piece, a trust anchor, if you like. And so it would seem to me we should take extraordinary measures, not necessarily expensive, but take extra measures to ensure that, in fact, that trust is well founded.

So having some rigor in how we look at security products is, I think, important. Independent evaluation is an important piece of that rigor. That is something different than the vendor claims. So where does one get this independent look? What is the most efficient way to get that independent look that in and of itself can be trusted by people who use these systems?

So whether it is a Common Criteria-based system that we use, whether it is some derivative that may be the result of an evolution of the process, I believe that we need to put honest faith in our security products through some independent specification evaluation process. It is too important just to sort of leave to the normal process.

Mr. PUTNAM. Mr. Gorrie.

Mr. GORRIE. There is two parts to this, as Mr. Fleming said. There is the independent evaluation portion of it, the Underwriters Laboratory, if you will. Should that be extended to the rest of the Federal Government? The Department of Homeland Security thinks it is. That is why they want us with them to do a review of the NIAP process, to see what that extensibility of the process is to the rest of the Federal Government.

Is it extensible to the rest of the civil population? No one forces a consumer to buy a lamp that has the Underwriters Laboratory stamp on the cord, and perhaps no one should be forced to buy a

piece of IT security equipment with a NIAP certificate associated with it.

There is the evaluation program itself, and then there is the regulatory and policy piece that goes along with it. And although I think the evaluation portion of it can go forward, because knowledge is power, the—how you instantiate that in regulation and in national policy is a different matter altogether.

Mr. PUTNAM. Thank you very much. And I want to thank all of our witnesses on this first panel and encourage you to stay and listen to the second panel, if your time and schedule allow.

With that, the committee will stand in recess for 2 minutes while the first panel is dismissed and the second panel is seated.

[Recess.]

Mr. PUTNAM. The committee will reconvene. Before we swear in the second panel, I did want to announce publicly that the executive session on SCADA, which was scheduled for tomorrow by the subcommittee, has been postponed thanks to Hurricane Isabel.

And with that, I would ask panel two to please rise and raise your right hands for swearing in.

[Witnesses sworn.]

Mr. PUTNAM. Note for the record that all of the witnesses responded in the affirmative, and we will move immediately to their testimony. Again, I would ask that you limit your remarks to 5 minutes, and your entire written statement will be submitted for the record.

Our first witness is David Thompson. Mr. Thompson directs the CygnaCom Security Evaluation Laboratory. He has led a team to support certification for the Air Force Scope Command's High Frequency voice and data communications system, and managed Public Key Infrastructure products at several Department of Energy National Labs. He led a team to write a Common Criteria security target for Red Hat Linux 5.1, and helped translate high-assurance criteria into Common Criteria protection profiles.

Previously Mr. Thompson evaluated the security of network and computing configurations for the space station and space shuttle, and assessed proposed uses of cryptography and distributed authentication at NASA. He was session chairman for the 1993 AIS Security Technology for Space Operations conference, and served on a board investigating a software configuration management failure in a space shuttle mission.

Welcome to the subcommittee. You are recognized, Mr. Thompson.

STATEMENT OF J. DAVID THOMPSON, DIRECTOR, SECURITY EVALUATION LABORATORY, CYGNACOM SOLUTIONS

Mr. THOMPSON. Thank you. I would like to thank the committee Chair and all of its members for their interest in this issue and their leadership.

The motivation for product testing that led to the creation of the Common Criteria came from the U.S. Government's certification and accreditation process for systems. Most systems included at least one computer with operating systems that needed a security functionality identified and assessed. Since operating systems are

quite complex and have many key security functions, considerable effort is required to do an appropriate security assessment.

As computers became more commodities, the notion of performing these difficult evaluations once and using the results in repeated CNAs took hold. In the early 1990's, as the expense of having products evaluated to different security criteria in different countries increased, Western governments began to seek a set of Common Criteria that they could endorse. We are still in the early stages of implementing the resulting Common Criteria. But the original government participants are still actively engaged, and additional governments are getting involved.

Industry also begun to see the value of a common security performance process. The CC defined seven sets of security assurances called evaluation assurance levels. EAL1 has the least assurance, and the EAL7 the most. The most commonly used assurance levels are EAL2 and EAL4. The EAL2 is an acceptable assurance level for most products, and EAL4 is often specified for products that are employed in the first line of defense, such as firewalls and operating systems.

Custom sets of CC assurances can also be chosen when one of the seven EALs is not precisely suitable. The result of a successful CC evaluation is a published security target that precisely documents the security functions that the product claims to meet and establishes in precise terms whether these claims are true, the security target to be used to determine the product's suitability for a particular use and to compare its security functionality with that of other products.

It is practically impossible to determine that a product of any complexity will be secure regardless of its configuration or that security will mean the same thing in all the situations in which the product is used. What CC testing does show is that to the specified level of assurance, the security functions the vendor claims the product has work as described, and that a coherent and mutually supported set of security functions is available.

Just because a product has been successfully evaluated under the Common Criteria does not mean that it has no vulnerabilities. Instead, it shows that the product is suitable for use as a component of a secure system. It is primarily focused on design and development process issues.

Although higher CC assurances, such as the EAL7, also can significantly reduce the possibility of bugs, the Common Criteria evaluation process has several strengths. It provides consumers with an independent and well-monitored assessment of vendor security claims. It provides a precise description of a product's security features that is readily comparable to those of other evaluated products. It assesses the ability of a product to be used to build secure systems. It demonstrates that at least one configuration of a product meets the claimed security requirements. It allows precise tailoring of the security criteria to the capabilities of products. It uncovers design flaws and sometimes software bugs. It focuses vendors on security issues. It constitutes the most rigorous and thorough independent product testing process commercially available. It provides international mutual recognition so that vendors have to pursue only one evaluation against a single criteria.

The Common Criteria evaluation process also has some drawbacks. It creates additional expense for product vendors. CC evaluation is applied to an exact version of a product in a precise hardware environment, making it sometimes hard to field a product that is strictly conformant. As consumer protection profiles evolve, as vendor products are revised, they must be reevaluated. The evaluation process is complex and time-consuming, which means it requires a lot of vendor resources, understanding and participation. Some of these are conflicting. For example, establishing the security of a product across a broad range of its configurations among many versions is more difficult and would further increase the expense.

While large vendors are more easily able to absorb the cost of an evaluation than smaller vendors, small vendors benefit more from an independent product assessment that makes it easier for customers to compare its products' security features to those of its better-known competitors.

The CCE offers a broad range of assurances and a corresponding broad range of costs. The EAL1 evaluation costs are in the low tens of thousands of dollars. EAL1 is adequate for many applications. Higher assurance has higher cost and is appropriate where the security risks are higher. The problem of eliminating security bugs from complex systems, such as those we read about regularly in the news, requires resources many orders of magnitude greater than those required for CC evaluations. There is little theory to support solutions to this problem, and it remains an art form.

The most productive approaches to bug elimination involve improved software engineering practices to prevent the introduction of bugs in the first place. Finding and fixing bugs in existing products is always more expensive.

The CC product evaluation process is a very effective tool when used in the right context. The international support and precise specification of security attributes minimizes the problems inherent in integrating diverse systems and components built in different countries and secure systems whose security attributes are well understood.

The Common Criteria evaluation, however, does not serve every purpose. The fact that attempts are made to apply it to situations for which it was not designed shows how great is the need for other kinds of security testing and the challenges facing the available security evaluation services.

Thank you for this opportunity to address you today.

Mr. PUTNAM. Thank you very much, Mr. Thompson. We are glad to have you.

[The prepared statement of Mr. Thompson follows:]

ANATOMY OF THE COMMON CRITERIA

Testimony before the House Government Reform Committee

Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census

September 17, 2003

Submitted by J. David Thompson

Director of the Security Evaluation Laboratory of CygnaCom Solutions, an Entrust Company

I. Brief description of Common Criteria

The motivation for a product testing capability is derived from the US military Certification and Accreditation (C&A) process for systems. Most systems include at least one computer, each employing an operating system that had to have its security functionality identified and assessed. Operating systems are complex and implement many key security functions, so considerable effort is required to do an appropriate security assessment of one. As computers became commodities, the notion of performing these difficult evaluations once and using the results in many C&As took hold.

The Orange book's set of five operating system criteria, the Rainbow Series of supporting methodology and interpretation documents written to supplement it, and the TPEP infrastructure, for the most part, accomplished this task. The Orange Book example was implemented in other countries, although in different and evolving ways. One of the evolutions was the European notion of a catalog of security assurance and functional requirements that could be used to specify security criteria for various types of IT components, not just operating systems.

In the early 1990's the Orange Book's narrow focus on operating systems became a problem, as was the expense to vendors of having to have their products evaluated to different security criteria in order to sell to the governments of different countries. A common criteria was sought that incorporated the best of the various existing programs and that all of the major Western governments could endorse. We are still in the early stages of implementing the resulting Common Criteria, but all the originators are still enthusiastic participants and additional governments are signing on to recognize CC conformance in their procurements and to help produce product evaluations.

Industry also sees the value of a common security conformance process and is using the CC's processes and flexible criteria to fit its purposes. The Trusted Computing Group (TCG), for example, has adopted the Common Criteria for specifying conformance to the security components of its Trusted Processing Module (TPM). The TPM standard specifies a chip that can be installed on a PC motherboard to support secure e-commerce, in effect turning PCs into smart cards. The TCG is a consortium of most major PC hardware and software manufacturers and TPM conformant chips are destined to be built into nearly all PCs. Some are already available today. The TCG specifies the CC evaluation process as the way for vendors to demonstrate the conformance of their chips to the security parts of the TPM standard.

The Banking Industry Technology Symposium (BITS) is in the process of rewriting its security criteria to be Common

Criteria conformant and replacing its ad hoc product testing facilities with testing in accredited CC testing labs. Their major criteria have been converted, and CC evaluations that include BITS conformance are underway.

II. Anatomy of the Common Criteria

The Common Criteria consists of two menus of security requirements -- one for security functionality (SFRs) and one for security assurance (SARs) -- and a process for using these components to evaluate products. The security assurance and security functional requirements are used to create documents that specify security criteria. The CC specifies two such documents: Security Targets and Protection Profiles.

A Protection Profile (PP) specifies security criteria for a class of products. A PP is intended to be written by a consumer or a group of consumers to specify the security requirements they want to see in a class of product they want to buy. The NSA has developed a large set of PPs for firewalls, operating systems, smart cards, and many other product types to specify security criteria for products for DoD use. BITS and the TCG are other examples of consumer groups producing PPs.

A Security Target (ST) specifies security criteria for a specific product. It also describes the security functions in the product and provides a rationale that the product's security functions meet the specified security functional requirements, among other things. It may also include an argument that its security criteria conform to one or more Protection Profiles. An ST evaluation confirms any PP conformance claims and the validity of correspondence arguments. A product evaluation confirms the underlying conformance of the product to the criteria in the ST.

The security assurance requirements identify a set of methodology that the CCTLS must execute satisfactorily in order to show that the product conforms to the security functionality in the security target. The security assurances used in an evaluation provide the consumer with a level of assurance that the product performs the security functions that the target says it will perform.

The CC defines sets of security assurances, called Evaluation Assurance Levels (EALs) that specify coherent groups of assurances roughly corresponding to those in the earlier Orange Book and ITSEC criteria. There are seven assurance levels defined by the CC, identified as EAL1 through EAL7. EAL1 specifies the least assurance and EAL7 the most. Other sets of assurances are viable, as well. The NSA has identified three sets of assurances associated with its Basic, Medium, and High Robustness levels. The authors of PPs and STs may, and often do, augment one of the EALs with additional SARs.

as meets their needs. These are often specified as EAL3+ or AL4+.

The most commonly used assurance levels are EAL2 and EAL4. EAL4 is often specified for products that are employed in the first line of defense, such as firewalls and operating systems, where untrusted users have full access to a feature-rich interface and the rewards of break-in are high. Lower assurances are appropriate for products with a lesser security role. Higher assurances are often required when information at different sensitivities must be separated or critical missions or assets are being protected.

The security assurances are grouped into the following classes:

- Configuration Management
- Delivery and Operation
- Development
- Guidance Documents
- Life Cycle Support
- Maintenance of Assurance
- Testing
- Vulnerability Assessment

The lower EALs do not include components from all of these classes, but it should be clear from this list that testing (trying to break into the product) is only one aspect of an evaluation, and only then in the context of an analysis of the design and implementation of the product. It is what we call white-box testing, requiring that the evaluator see inside the box and understand its functionality before running tests. Vendor cooperation is required. The alternative is black box testing, where nothing is known about the inside and the interfaces are tested against their specification. Black box testing is generally cheaper but provides significantly less assurance.

The result of a successful CC evaluation is a published Security Target that precisely documents the security functions that the product claims to meet and provides a precise expression of the assurance that has been applied to confirming the claims are true. The ST can be used to determine the product's suitability to a particular security use. It can also be used to compare the security functionality of competing products in a way that vendor marketing information makes difficult, if not impossible.

It is theoretically and practically impossible to determine that a product of any complexity will be secure regardless of its configuration, or that security will mean the same thing in all the situations in which the product will be used. What CC testing does show is that the security functions the vendor claims the product to have work as described and that a coherent and mutually supportive set of security functions is available.

III. Misconceptions

Two common misconceptions are worth addressing. The first arises from the fact that many products that have been successfully evaluated are later found to have vulnerabilities. This leads to the conclusion that the evaluation process must be somehow flawed. In truth, the evaluation accurately demonstrated that the product can be used to implement secure systems in a specific configuration. The problem arises when vendors or users of the product do not use it in a secure way, choosing instead to misuse the security features to maximize utility or ease of use at the expense of security. The NSA, among others, publishes guidance for securely configuring critical widely used products on its website, at www.nsa.gov/snac.

Another misconception is that a product that has been evaluated should be free of bugs. The CC evaluation process is primarily focused on design and development process issues, not finding bugs. Higher CC assurances, such as those that comprise EAL7, do reduce the possibility of bugs significantly since they require that product design and development proceed in parallel with the evaluation and that important design documents be expressed mathematically and proven to correspond to each other. Further precise documentation and thorough correspondence between documentation layers is provided down to the code level, and anomalies of compiler and processor design are taken into consideration. But an EAL7 evaluation is only practical on small programs (less than 10,000 lines of code) when several million dollars are available. Developing perfect code, to perform security or other even moderately complex functions, is exceedingly difficult and out of the price range of even an EAL7 evaluation by several orders of magnitude.

NSA is developing protection profiles that specify small operating system kernels with simple security functions for use in special purpose high-risk applications that can be evaluated to a high assurance level. The most complete example is the Partitioning Kernel PP for Real Time operating systems. These systems must isolate data at multiple security levels (e.g., Top Secret and Unclassified). That PP specifies EAL7+ and several vendors are pursuing the development of such systems and their evaluation against it.

IV. Strengths

The Common Criteria evaluation process has several strengths, which are listed below:

It provides consumers with an independent and well-monitored assessment of vendor security claims. These claims are often difficult to determine from marketing literature that touts security features, or even from independent reviews that compare products.

It provides a precise expression of a product's security features that is readily comparable to those of other evaluated products. This description is similar to that used in legal documents, with carefully defined terms. Consequently, it allows comparisons between different expressions to be made more accurately.

It assesses the ability of a product to be used to build secure systems. It clearly identifies the security functions and the limits of their implementation.

It demonstrates that at least one configuration of a product meets the claimed security requirements. As part of the evaluation, the vendor must specify a secure configuration of the product. That configuration then becomes the basis for the vulnerability analysis and the vendor and evaluator testing.

It allows precise tailoring of the criteria to the security capabilities of products. The flexible nature of the CC's menu of security functional requirements allows the specification of nearly every security function and its customization to the precise method implemented. The ability to augment the CC requirements with modified or completely new requirements allows the complete specification of any security function any product might have.

It uncovers design flaws and, sometimes, software bugs. The CC process is best at uncovering design flaws. In some cases, the perspective of a CC evaluation often leads vendors to see security design flaws that they didn't recognize as flaws before. Sometimes, it also uncovers bugs.

It focuses vendors on security issues. Some vendors do not spend much time worrying about security. A CC evaluation directs their energy into security and makes them defend their security designs to an independent third party.

It constitutes the most rigorous and thorough independent product testing process commercially available. There are other independent testing processes that are cheaper and less intensive, but the CC is the most fundamental. Without it we would be much less able to select the right products to build secure systems or to understand the risk remaining in those systems.

It provides International Mutual Recognition, so that vendors only have to pursue one evaluation against a single criterion. This is an important advantage of the CC over its predecessors. It provides a larger market over which to amortize evaluation costs.

V. Weaknesses

The Common Criteria evaluation process also has some weaknesses, which are described below.

It creates an additional expense for product vendors. CCTL fees range from \$30K or less for an EAL1 evaluation to over a \$1M for an EAL7 evaluation. The cost to the vendor to support the evaluation may be in the same range as the CCTL fees, effectively doubling these estimates. The necessary EAL is determined by customer requirements and by competition. The evaluation specifies a precise version of the product and a precise hardware environment. Other versions and hardware platforms are not strictly evaluated. This is due to the fact that letting those parameters vary makes it nearly impossible to reach a meaningful conclusion about the security of the product. Some consumers require strict conformance. The risk assessment part of the C&A process must deal with residual risk imposed by deviations from the evaluated configuration.

As products protection profiles evolve, products must be re-evaluated over their life cycle.

Because the CC evaluation process is complex and time-consuming, it requires a lot of vendor understanding and participation.

VI. Conclusion

The CC product evaluation process is a very effective tool for a very important purpose. It is critical for countering the growing threat arising from the convergence of global software development and international terrorism. Its wide international support and precise specification of security attributes minimizes the problems inherent in integrating systems and components built in different countries and services into effective and secure systems—systems whose security attributes are well understood. It does not, however, serve every purpose. The fact that attempts are made to apply it to situations for which it was not designed shows how great the need is for other kinds of security testing and how many challenges face the available security evaluation services.

Mr. PUTNAM. Our next witness is Mary Anne Davidson. Ms. Davidson is the chief security officer at Oracle Corp., where she has been for the last 14 years. As Oracle's CSO, she is responsible for Oracle product security, corporate infrastructure security and security policies, as well as security evaluations, assessments and incident handling.

Ms. Davidson also represents Oracle on the Board of Directors of the Information Technology Information Security Analysis Center, and is on the editorial review board of the Secure Business Quarterly.

Prior to joining Oracle in 1988, Ms. Davidson served as a commissioned officer in U.S. Navy Civil Engineer Corps, during which time she was awarded the Navy Achievement Medal. She has a BSME from the University of Virginia, and an MBA from Wharton at the University of Pennsylvania.

We always appreciate your interaction with this subcommittee and your direct and candid remarks. Welcome. You are recognized.

STATEMENT OF MARY ANN DAVIDSON, CHIEF SECURITY OFFICER, SERVER TECHNOLOGY PLATFORMS, ORACLE

Ms. DAVIDSON. Mr. Chairman, Ranking Member Clay, on behalf of Oracle, I appreciate the opportunity to be here today to offer Oracle's perspective on the Common Criteria.

Oracle is uniquely qualified to comment on information assurance policies. We have spent more than 25 years building information management systems for customers that I affectionately call the professional paranoids, which include U.S. intelligence agencies and the Defense Department.

To gain and maintain the business of the most security-conscious customers on the planet, we have made extraordinary investment in information assurance and have 17 independent security evaluations to show for it, with 4 more in process.

The collective impact of Code Red, Blaster and SoBig to our economy, which amounts to billions of dollars in repairs and down time, have worked to send a sobering message: It is long past time for the entire Federal Government to get serious about information assurance. The benefits go beyond secure Federal information systems. A strong Federal information assurance policy has a potential to change the entire software industry for the better. Let me tell you there is no vendor when faced with this requirement who will build two versions of software, one that is strong, robust and well-engineered, and a buggy, crummy version for the commercial sector.

Fortunately, some Federal agencies are listening. NSTISSP 11 and DOD Directive 8500.1 draw a constructive, clear prosecutry line in the sand. The question is not whether NSTISSP 11 makes sense or not. We have had that debate, and it is over.

The NSTISSP 11, with the linkage to the Common Criteria, the de facto worldwide evaluation standard, is making a positive, constructive difference in software development. The Common Criteria has three key benefits for vendors who do evaluations: You have more secure products. Evaluators find security vulnerabilities which must be remedied prior to receive a certificate. There has been a lot of discussion about the cost of evaluations, but I have

done the analysis, and if we find or prevent even one significant security flaw in our products going through the evaluation, it more than pays for the cost of the evaluation, even at the highest assurance levels that are viable for commercial software, which says nothing about the expense to customers that we prevent by getting it right the first time.

Second, a more secure development process. Evaluations actually force you to have a secure development process throughout the entire development process. Security can't be something that is thrown on in the last 2 weeks of a cycle and has to be baked into the development process. That is what the evaluators are looking at.

Third, and probably most important, a strong culture of security. If you do evaluations as part of development, then security becomes baked into your corporate culture. That is the biggest problem that we have in the industry: Security always seems to be someone else's job. At Oracle I tell our developers, you are personally responsible and accountable for every single line of code you write.

Since NSTISSP 11 has gone into effect, we have seen very positive developments. More firms are doing evaluations. Firms, including Oracle, are sponsoring open-source evaluations. Many other industries are looking at certification efforts along the same lines as the Common Criteria. This has been successful because industry believes the Federal Government is serious this time, and that is a major victory. And thanks goes to people within DOD, the Intel Community, and Congress, who are making an effort to make the process work.

So what can we do to make this process work better? You can hold the line by maintaining an eternal but pragmatic vigilance through the no-waivers policy. I said a year ago that it was time for the government to chirp or get off the twig on information assurance, and there has been a lot of chirping going on.

But there are still those who want to get off the twig by getting a waiver or seeking opt-outs. It sends a bad message to the marketplace to say that NSTISSP 11 does not apply to us. It really needs to apply to intelligence across the board.

NSTISSP 11 should be extended beyond traditional national security systems, and specifically, I think DHS should look at applying this to their own systems. Clearly they have a mission of national security.

NSTISSP 11 shouldn't be allowed to—protection profiles should not be agency-specific wish lists. I think vendors are willing to do an evaluation against a common protection profile, but they are not willing to do three of them per each special agency.

Country independence of laboratories should be maintained. We do our evaluations in the United Kingdom, because the cost is lower and the expertise is actually far higher than we have found in the United States for our particular product set. We still get resistance to foreign evaluations, and this is ridiculous. We are very happy to support U.S. labs as a competitive alternative, but competence knows no national boundaries.

A couple of final points. There are three things that the government can do to foster better security beyond evaluations. We know that it does not provide a silver bullet or perfect products. The Fed-

eral Government should require that products have a default setting that is secure out of the box. I think NIST can do a lot of work here. This would also provide a lot of immunity to a number of viruses and worms, because more systems would be locked down by default. It would lower the cost of operations for the government and other customers.

The government should invest in cybersecurity research. Quite honestly, the reason vendors cannot find more faults in the products in development is because the tools do not exist to do so, and the venture capital community will not fund it because there is no way to make money on it. If we can stomp out smallpox through investing in medical research, we can certainly get rid of buffer overflows. It is just code.

Finally, industry can do more to improve the security profession. I fully support an alternative to Common Criteria evaluations, for example, for consumer products where it is perhaps inappropriate to do a Common Criteria evaluation. And an example would be the Underwriters Lab. Most products are just designed to be secure. And Cuisinarts are designed, for example, that you can't lose your fingers by sticking them in while the blades are whirring. They are just secure. Consumers don't have to do something special to make them operate securely.

NSTISSP 11, DOD 8500.1, and the national strategy are welcome developments because they are moving the debate to the expectation that everything will be secure. I believe we have turned a corner, but it took 10 years and numerous sobering events to get us here. It will take continued vigilance and continued leadership here in Congress to keep us on this road.

Thank you again, Mr. Chairman, for the opportunity to testify today.

Mr. PUTNAM. Thank you very much, Ms. Davidson.

[The prepared statement of Ms. Davidson follows:]



Statement of

Mary Ann Davidson
Chief Security Officer
Oracle Corporation

Before the

Subcommittee on Technology, Information Policy,
Intergovernmental Relations and the Census
Committee on Government Reform

US House of Representatives

17 September 2003

Mr. Chairman, Ranking Member Clay, my name is Mary Ann Davidson, and I am the Chief Security Officer of Oracle Corporation. On behalf of Oracle, I appreciate the opportunity to be here today to offer Oracle's perspective on information security, and specifically, the Common Criteria. This is a critically important topic, especially given events over the past month.

Oracle is the world leader in enterprise software, and is uniquely qualified to comment on information assurance policies. We have spent 25 years building information management systems for customers that I affectionately call 'the paranoids,' which include US intelligence agencies and the Department of Defense. To gain and maintain the business of the most security-conscious customers on the planet, Oracle has made an extraordinary investment in information assurance, and we have 17 independent security evaluations to show for it. The basis of our marketing campaign "Unbreakable" is this long-term commitment to information assurance.

We made this investment in security for one simple reason: Our customers asked for it. They asked for it, and they meant it. Up until recently, we have witnessed what I could call a merry-go-round on information assurance within the federal government. Despite more than ten years of well-intentioned efforts by federal agencies to ask software vendors to have their products independently evaluated, vendors simply refused to do them because, in the end, they counted on the federal government to not follow through on its own request. Meanwhile, the federal government would refuse to get serious about evaluations because not enough vendors did them. A lazy vendor too often was just a weasel-willed procurement officer away from cheating on evaluations.

The collective impact of Code Red, Blaster and Sobe to our economy, which amounts to billions of dollars in repairs and downtime, have worked to send all of us a sobering message: It's time to get off the merry-go-round once and for all. The benefits go beyond more than just secure federal information systems. A strong federal information assurance policy has the potential to change the entire software industry for the better.

Fortunately, some federal agencies are listening. The message has come largely from the policy directive mentioned numerous times already today: NSTISSP #11. This policy, as well as several enforcement components, most notably Department of Defense Directive 8500.1, drew a constructive, clear, pro-security line in the sand. Simply put, for national security systems, an agency can only purchase commercial software that has been independently evaluated under the Common Criteria or the Federal Information Processing Standards (FIPS) Cryptomodule Validation Program (CMVP).

Mr. Chairman, the question before us is not whether NSTISSP #11 makes sense or not. We've had that debate. It's over. NSTISSP #11 is already making a positive, constructive difference in software development. Instead, the question before us is how to make this policy work as effectively as possible and within as many federal agencies as possible.

As we all know, the success of NSTISSP #11 is linked in part to vendors participating in the Common Criteria, the de facto worldwide evaluation standard, which has the added benefit of mutual recognition by many countries, including the US, the UK, Germany, Canada, France, Australia and New Zealand. Vendors complete one security evaluation that is valid in many countries. Consumers of the software have assurance that the vendor is not blowing smoke, because it is someone other than the vendor validating security claims. (Let's face it, all vendors claim they are secure, even the ones who issue security patches for their products every 2 ½ days.) By being linked to the Common Criteria, NSTISSP #11 has three key results:

- First, more secure products. Evaluators find security vulnerabilities, which must be fixed. No fix, no evaluation certificate, no exceptions.
- Second, a more secure development process. Evaluations actually test the process more than the product itself. Product security architecture, functional, design, and test specifications are reviewed, and a secure development process has to be repeatable. Security can only be built in from inception, not "bolted on" after the fact.
- Third, a stronger culture of security. Instituting evaluations as part of software development, and then repeating them over and over changes the corporate culture. Security becomes part of the corporate "DNA," woven into the fabric of the organization. This is the biggest long-term benefit of security evaluations, because over time, it becomes an industry culture.

It's been 14 months since NSTISSP #11 has gone into effect, and we have seen several very positive developments. First, a number of firms, including several of our competitors, are getting their products evaluated under FIPS or the Common Criteria, and some for the first time. Second, we're seeing firms, including Oracle, financing evaluations of open source products, which will work to dispense some of the myths surrounding the so-called inherent vulnerabilities in open source operating systems. Third, several industry organizations, such as the financial services industry, are coming together to make security a purchasing criterion industry wide and are using NSTISSP #11 as a model.

We're seeing all of this because the initial impression from an industry perspective is that the federal government means business this time. That, in and of itself, is a major victory and credit goes to the people within the Defense Department and intelligence agencies, as well as Congress, who are making a concerted effort to make this process work. The fundamental question for them, you and other policy leaders, and all of us in the security world is how can we continue to make this process work even better. Let me provide a few suggestions:

First, maintain eternal, but pragmatic, vigilance. Enforcement of NSTISSP #11 should be consistent, with no waivers. Fortunately, the enforcement process was set up to discourage waivers, and shifts waiver authority from procurement officers to the Committee on National Security Systems within the NSA, which was the entity that first developed NSTISSP #11. The merry-go-round that I referenced earlier was driven by a decade of procurement “dodge and weave” via waivers. Last year, in testimony before the House Readiness Subcommittee, I said that when it comes to information assurance, it was time for the federal government to “chirp or get off the twig.” So far, so good. The cat hasn’t eaten the canary yet.

Don’t get me wrong. Several agencies or sub-agencies wouldn’t mind getting off the twig, and opt out of NSTISSP #11. A general sentiment of “NSTISSP #11 does not apply to us,” especially when so many components used by intelligence agencies are standard commercial software and hardware products, sends a terrible message to the marketplace and negates the intent of NSTISSP #11. What could be more central to national security than intelligence?

Second, the federal government should extend NSTISSP #11 beyond traditional national security systems. The creation of the Department of Homeland Security, coupled with the increasing importance of federal government information systems to maintain the effective administration of health care, social welfare and law enforcement requires the entire federal government to make security a factor in software buying decisions. Earlier this year, the President’s National Strategy to Secure Cyberspace recommended a study on the application of NSTISSP #11 across the entire federal government. I encourage this subcommittee to check on the status of this study, work to see it is completed and ensure that representatives of industry have the opportunity to make recommendations. Last year, while debating the creation of the Department of Homeland Security, several Senators, including the former Chair of the Senate Government Affairs Committee, agreed that the new Department should examine how NSTISSP #11 can be implemented in their own procurement policies. I also would encourage this subcommittee to call on the Department of Homeland Security to institute such a policy.

Third, to make NSTISSP #11 work across the federal space, protection profiles should not become agency-centric wish lists. As you well know, Mr. Chairman, evaluations are not cheap. In order to ensure vendors get maximum return on their evaluation investment, evaluated products need to be commercially viable. We have seen protection profiles specifying requirements that no commercial product can meet or for which there is no commercial requirement. If a federal agency wants something new and different, the correct vehicle should be something other than a protection profile. Oracle supports such industry-government dialogues, but the corollary to ‘if you build it, we will buy it’ needs to be ‘can this be built, or is there a better way of solving the problem?’ This often puts us at odds with security purists who design things that are so secure they are undeployable. Fortezza, for example, was never widely-deployed and was given a rousing thumbs’ down by the commercial marketplace, as well as many DoD customers

who were supposed to use it. The marketplace ultimately delivered commercially viable encryption, and FIPS evaluations can ensure that the implementations are done correctly.

Similarly, I would recommend that if a federal protection profile effectively requires a Common Criteria evaluation level above EAL4, or imposes a requirement that is not commercially viable, the federal government should pony up the money for the evaluation. Again, the goal here is to make sure that vendors get the investment return. The key commercial advantage of the Common Criteria is mutual recognition -- one evaluation works typically for all nations that are signatories of the Common Criteria Mutual Recognition Agreement. Specifically, at an EAL4 assurance level, mutual recognition applies; at higher levels or at hybrid levels (EAL4 plus some other requirements), mutual recognition is void. Most vendors are willing to do one evaluation — which can run as much as \$1 million per product — as the cost of doing business with the federal government, but a situation where agency specific protection profiles result in “three evaluations per product” is cost-prohibitive. We’ve been down this road before. Oracle has evaluated products against the US Trusted Computer Systems Evaluation Criteria (TCSEC or “Orange Book”), the UK Information Technology Security Evaluation Criteria (ITSEC), and the Russian Federation Criteria to meet per-country evaluation requirements. We would propose that any procurements requiring assurance levels higher than EAL4 or which void mutual recognition must foot the bill for the evaluation as part of the procurement, including vendor personnel costs. EAL4 assurance is attainable by commercial software products; anything higher is generally non-commercially viable and the procurement officials should expect to pay for a custom solution and a custom evaluation, and budget accordingly.

I know the Chairman has expressed previously his concerns about the cost of security evaluations. Streamlining the protection profile process to prevent unnecessary evaluations can reduce costs, and maintain the commercial viability of the product, without compromising the security benefits of the Common Criteria.

Fourth, country independence of laboratories should be maintained. For business reasons, Oracle’s security evaluation group is headquartered in the United Kingdom and we almost always use evaluation laboratories in the United Kingdom. The process has been wildly successful, with seventeen evaluation certificates under our belt. NSTISSP #11 avoids a nation-centric approach to labs, especially given the obvious benefits to vendors of the mutual recognition provisions of the Common Criteria. That said, we still experience resistance from procurement officers to “foreign” evaluations. This is nonsense. The Common Criteria supports mutual recognition, and a non-US laboratory can become certified to perform (US) FIPS-140 certifications. Also, the Common Criteria process works from objective evaluation standards, and labs, regardless of location, must meet objective standards to become a Common Criteria evaluation facility. Thus, there is slim reason to suspect foreign labs.

We cannot return to the days of country-specific evaluations, which will be the end result if procurement resistance to foreign evaluations spirals down into a policy requirement.

We certainly support efforts to make US labs a competitive alternative to foreign-based labs, which would help reduce evaluation costs in the long run. Competence knows no national boundaries; neither does incompetence.

Fifth, the federal government should establish a clearinghouse on evaluation

product information. There are already several good web sites to help both vendors and their federal customers understand Common Criteria, FIPS, and NSTISSP #11. However, we find that what many of our customers need is a one stop, 'go to' site in order to validate vendor security claims and compare them to the evaluation results themselves. It would be useful for a procurement officer to be able to see all evaluations of any type, for a single vendor, at a single glance, from a single location, whether FIPS-140 or Common Criteria, whether evaluated here or abroad. This empowers them to make apples to apples comparisons. For example, two database vendors can both receive an EAL4 certification, even though one database vendor made two functionality claims in a security target, while the other database vendor made forty security claims. A clearinghouse would enable a procurement officer to perform security target 'scorecarding' and facilitate this and other types of comparisons.

Mr. Chairman, there are no security magic bullets, and certainly, NSTISSP#11 is not meant to be one. What it does provide is assurance that the allegedly secure gun does what its seller says it does, without misfiring and killing the user. The stakes have never been higher for information security, and especially information assurance.

The fundamental lesson of the last 14 months is clear: if the federal government acts like a buyer concerned about the inherent security of its software, its sheer market presence alone will change the behavior of vendors for the better. That said, are there other ways, other than NSTISSP #11, that can accomplish the same purpose? We believe one measure worth considering is for the federal government to insist that the commercial software it buys is either defaulted to a secure setting right out of the box, or made easy for the customer to change security settings. OMB, working in conjunction with the National Institute of Standards and Technology (NIST) and private industry, can specify what is the appropriate default security setting for the software it buys.

We also can't emphasize enough the value of independent research. Assurance is not only evaluations, and even with a good development process, "to err is human." A developer can check 20 of 21 conditions, and if failure to check the 21st causes a buffer overflow, the system is still potentially vulnerable. Hackers only need to find one error, while developers have to anticipate and close every one. It's an uneven battle. The federal government, working with academia, has the ability to jump start research that moves the security ball down the field. One area that deserves attention, especially as more and more US firms partner with foreign countries on software development, is research on effective tools that can scan software and pinpoint irregularities or backdoors in the code. While one would think there would be a market for such software, the research and development involved is seen as cost prohibitive for the private sector.

Fortunately, Congress last year passed the Cyber Security Research and Development Act, which authorizes funds for projects like code-scanning tools.

If the medical community could eradicate smallpox with a strong investment in research, we should be able to eradicate buffer overflows.

Finally, industry can do more voluntarily, and in conjunction with academia, to establish professional standards for security officers – professional standards that would evolve as security practices and software development evolved. The software industry also should come together voluntarily to establish the software equivalent of the Underwriters Laboratory. Security evaluations under the Common Criteria are not necessarily cost effective for many forms of consumer software, especially given a price of hundreds of thousands of dollars per evaluation for large, complex products at relatively high levels of assurance. Again, the fundamental goal is to make all commercial software secure by default. To get there, the federal government should work with private industry to establish a consumer software equivalent of the Underwriters Laboratory (UL). Thanks to the UL, most consumer products are generally difficult to operate in an insecure fashion. For example, Cuisinarts are designed so that you can't lose a finger while the blades are whirling. We don't expect the consumer to do anything special to operate Cuisinarts securely; they just are secure.

Far too much commercial software today is built without attention to information assurance principles, leaving many of our national cyber-assets easily vulnerable to an attack. With advanced hacking tools easier to obtain, and the global economy's increasing dependence on web-based platforms to perform everything from financial management to intelligence gathering and analysis, we can no longer patch our way to better security. Instead, we need to move toward a different approach, culture if you will, in which all software sold in the commercial marketplace is secure by default.

NSTISSP #11, DoD 8500.1, and the President's National Strategy to Secure Cyberspace are all welcome developments, because a common theme in all of these documents is the ability of the federal government to use its unique resources on the side of those who adhere to the 'secure by default' culture. I believe we have turned a corner, but it took ten years, and numerous sobering events to get us there. It will take continued vigilance, and continued leadership here in Congress and the Administration to keep us on this road.

Thank you again, Mr. Chairman, for the opportunity to testify today.

Mr. PUTNAM. Mr. Clay, if I heard her correctly, she said that she tells her developers that they are personally responsible for every line of code that they write. It is a good thing nobody holds us to that standard on the U.S. Code.

Our next witness is Mr. Klaus. Christopher W. Klaus is the founder and chief technology officer of Internet Security Systems, Inc., a leading global provider of information protection solutions that secure IT infrastructure and defend key online assets from attack and misuse. Prior to founding Internet Security Systems, Mr. Klaus developed the Internet Scanner, the first vulnerability scanner, while attending the Georgia Institute of Technology.

Mr. Klaus was honored in MIT's magazine, Technology Review. In addition, he received the award for Ernst & Young's Entrepreneur of the Year in 1999, in the category of Internet products and services.

Welcome to the subcommittee. You are recognized.

**STATEMENT OF CHRISTOPHER W. KLAUS, CHIEF
TECHNOLOGY OFFICER, INTERNET SECURITY SYSTEMS, INC.**

Mr. KLAUS. Thank you, Mr. Chairman. It is an honor to testify today. And I am representing Internet Security Systems from a small company's point of view that builds the security products, and we are in the process of going through the Common Criteria and NIAP certification and wanted to share some of our experiences as a company going through it, and what are some of the benefits and some of the failures of the process that we see today.

We do believe the overall goal and the intent of the Common Criteria and going through NIAP certification is a positive goal, but we see that there is at least three areas of major improvement that need to happen. And if they do not get addressed, we believe that following this path of requiring the government to follow the guidelines of NIAP certification actually makes the government less secure. And we will go through these three reasons and talk to why do they make both the government and others that follow the certification less secure.

No. 1 would be the accuracy. The current different levels of evaluation do not reflect whether the security product is actually more accurate in protecting against vulnerabilities and exposures. To take a step back, let me explain two goals within security, so you understand what we are measuring.

There is two major goals in security. One is to allow good people into the network, or into an operating system, into an application. And what we typically think of good guys in technology is like your user name and password that allows you to get into the system. There is biometrics, fingerprinting, VPN, virtual private networks. All of those technologies are great for—help certify the right people get into the system.

And one of the problems, though, is it assumes that the infrastructure stops the bad guys out. So the second goal of security is keeping bad guys out. The problem we find is that the assumption that the infrastructure keeps the bad guys out is false. We know there is everyday bugs in the code. These bugs lead to vulnerabilities that then allow intruders, worms, viruses to leverage that.

So on the second goal of keeping the bad guys out, that is a major area of measurement. And one of the things that we track very closely as a company that produces security products, we are tracking over 200 plus vulnerabilities every 3 months, every quarter, as we measure that, and what is interesting about this is as we go through this process, products that are less accurate in finding those vulnerabilities have the same certification as the companies that have much more accurate products. And if you likened it to antivirus, which most people are familiar with antivirus software, if only 10 percent of the vulnerabilities were—or 10 percent of the viruses were found with one product, and 99 percent were found with another product, today they would be measured equal in terms of the certification level. And that is one of the major reasons why government agencies that believe they are getting a more robust product may end up—just because they are purchasing a higher level certified product may actually end up with a less robust and less accurate product.

The next major area is speed. The current evaluation process is extremely slow and bureaucratic. It can take over a year to become certified. By the time it does become certified, it is outdated and behind the latest version of protection. The commercial sector could apply the latest version, and while the government would lag behind in security, in the race against cybercrimes threats, all organizations need to apply the current, most up-to-date security protection products.

I just would add that there is over 40 IDs or intrusion detection companies in the process today, but only two of them have actually been certified. So we have a long way to go before all of them have gone through this process.

The issue, though, also with security products, we are in a much different stage in the technology industry where we are rapidly evolving the technology to keep up with the cyberthreats. A lot of older technologies, like operating systems, data bases, Web servers, those technologies have been around longer, more mature, have a much longer development cycle. So in many cases the larger the application, and the larger the deployment cycle, the more likely you can keep in pace with the certification. In the security circles it is at a much faster rate.

And finally, the cost part of this is that the current evaluation process is extremely burdensome and costly for security vendors to follow. And after following the process, the expense does not—for us has not resulted in any security improvement. It has not found any buffer overflows. It has not found anything that many of the hackers and worms and viruses take advantage of. So, therefore, many of the resources and capital that we are spending on this, if it was doing something to make our products a lot better, and more protected, and more robust, and more accurate, we would be in favor of this.

So those three things, accuracy, speed and cost, are critical to improving, to make this thing worthwhile.

Mr. PUTNAM. Thank you very much.

[The prepared statement of Mr. Klaus follows:]



ISS Testimony Regarding Common Criteria practice

Submitted to the House Subcommittee on Technology, Information Policy,
Intergovernmental Relations and the Census

by
Chris Klaus, Founder and CTO, Internet Security Systems, Inc.

The overall goal and intent of Common Criteria and NIAP certification of helping the government select the proper levels of security products is good, but the actual process and mechanics fail in its mission to properly evaluate and assess the robustness of security products. It fails in three major ways: accuracy, speed, and cost.

Accuracy: The current different levels of evaluation do not reflect whether the security product is actually more accurate in protecting against vulnerabilities and exposures. A government agency could falsely believe they have better protection with a higher level of certified products, but in reality, have less robust and accurate security products.

Speed: The current evaluation process is extremely bureaucratic and slow. It can take over a year before a product becomes certified. By the time a product becomes certified, it is outdated and behind the latest version of protection. The commercial sector could apply the latest version, while the government would lag behind in security. In the race against the cyber-crime threats, all organizations need to apply current security protection products.

Cost: The current evaluation process is extremely burdensome and costly for the security vendor to follow, and after following the process, the expense does not result in any security improvements in the products for the government. The resources and financial capital is better spent on making more robust and accurate security products. The evaluation process should reflect these improvements.

Because the current evaluation process fails in these three major areas, the government will actually become less secure if it follows the Common Criteria guidelines. The criteria and certification process needs to be dramatically revamped and overhauled with much stronger participation and input between government and the commercial sector for the certification process to improve.

Here are some additional details on some of the key points about the Common Criteria:

Time-Line: For the minimum required level for government certification, EAL-2, the process requires a time-line longer than 7 months to complete. By the time a commercial product has passed government certification, it will be significantly out-dated. With the race against viruses, worms, and hackers, the certification process gives the cyberspace threats an advantage. The certification process needs to be designed to be more nimble and efficient, so that security products can quickly get certified and deployed rapidly within the government agencies to fight cyber threats.

GOTS vs. COTS: While the government is requiring a burdensome certification process against commercial-off-the-shelf (COTS) security products, government-off-the-shelf (GOTS) products do not need to meet the same level of standards for certification. Because Common Criteria is not being applied equally to both GOTS and COTS, many of the GOTS security products will avoid being held to the same standard of analysis and comparison. This makes it difficult for government agencies to determine the quality and make decisions on whether they are applying the best security solution to their network infrastructure. In head to head bake-offs, the COTS products are more comprehensive and have a more extensive level of protection than internally funded GOTS products. As the certification process evolves, it is important to apply the same process to both GOTS and COTS security products to ensure that all security products are being fairly and equally measured against the same standards and process.

Commercial Certifications: There are many commercial certifications that measure and certify security products for their quality of security content and protection capability. Typically, the commercial certification focus has been around firewalls, anti-virus, and intrusion detection/protection systems. They are designed and evolve with the input from the commercial industry so that they fairly quantify the security levels while minimizing the burden and time requirements placed on the security vendors. Leveraging the knowledge and lessons that commercial certification organizations have learned, it would be beneficial to government if this expertise were applied to the Common Criteria requirements.

In conclusion, Internet Security Systems believes the original goal of Common Criteria for helping governments to select appropriate levels of security products was good. From our business experience of going through the implementation and the mechanics of the certification process, it has clearly failed on actually measuring appropriate security levels. It has failed to keep the certification process within a time frame that is meaningful for cyber-threats. Common Criteria moves both the security industry and government agencies relying on the certification in the wrong direction. There are

commercial certification organizations that have developed very valuable certification criteria, while keeping the overall process, resources, and time requirements to a minimum. The government would be best to learn and leverage what commercial certification organizations have developed and get input and involvement from the vendors and non-government groups (like research and universities) to participate to help optimize the process and maximize the overall security goals for the good of national security.

Mr. PUTNAM. Our next witness, our last witness on this panel, is Eugene Spafford. Dr. Spafford currently serves as the director of the Center for Education and Research in Information Assurance and Security at Purdue University, a position he has held for 5 years.

He has written and spoken extensively on the topic of information security. His research focuses on the prevention, detection, and remediation of information security failures and misuse, including fault tolerance, software testing and debugging, and security policies.

He holds a Ph.D. in information computer science from the Georgia Institute of Technology.

We are delighted to have this level of expertise on the panel. And you are recognized for 5 minutes.

STATEMENT OF EUGENE H. SPAFFORD, PROFESSOR AND DIRECTOR, CENTER FOR EDUCATION AND RESEARCH IN INFORMATION ASSURANCE AND SECURITY, PURDUE UNIVERSITY

Mr. SPAFFORD. Thank you, Mr. Chairman. And thank you also, Ranking Member Clay and members of the committee.

The question posed to us for this hearing was can the Common Criteria ensure security for the Federal Government? And my answer to that is very definitely not. It will not.

And that is not to say that the Common Criteria is not a valuable instrument. The many thousands of man-years of effort by experts around the world putting it together has resulted in a procedure and set of documents that have great value as guidance for those building systems and for a means to compare systems as to their level of quality. However, it does not actually address the problem of ensuring that the government systems or any systems that possess the certification are themselves secure. It is in some sense, if I may use the analogy, similar to wanting to be sure that your house will not burn down and believing that the Underwriters Laboratory seal on the cord of your toaster will ensure that. It is not the case. What it does do is it gives you a small added measure that one item involved is less likely to cause you damage, but it is certainly no guarantee for the whole enterprise.

We can see that with an example that has been cited by many others. If we look at the Windows 2000 operating system, it is certified at the highest currently available level available under the Common Criteria, and yet it was a target. It was vulnerable to the Blaster worm, the Natchi Worm, dozens if not hundreds of current viruses, and has had nearly 100 patches issued for it—those are security patches, not functionality patches—since it was released. And that is something that is certified at the highest level.

There are other examples. I have given a detailed list in my written testimony as to why I do not believe that the Common Criteria is going to give us the level of security that we want. And in the very limited time available here, what I am going to do is give a different approach to this, and I am going to do it by analogy.

Let's take that toaster example that I was talking about. Suppose that you were the vendor of that toaster, and you wanted to compete on the market and decided that an evaluation was some-

thing that would give you a competitive advantage. So you submit it to a consumer testing lab. However, when you submit it to the testing lab, you submit it without a cord, and you tell the testing agency that you want to submit it as a bread storage device.

Well, the agency is required to test it against the requirements that you gave them. So they will test it as a bread storage device. They will go against the checklist for all the devices in the kitchen, and they will discover there are no radioactive materials or explosives embedded in it, and that, in fact, it does meet all of the documentation that you provided, it was built by the engineers in the appropriate way, and it does store slices of bread. So they give you their highest level of certification.

You then turn around and put it in the marketplace with a cord and with a consumer option to include a speaker phone, because while you are making toast, you want to call your neighbors and tell them to come over and have some of the toast. The problem there is that about every tenth piece of toast that you put in burns and possibly starts a fire. What is more, the speaker phone is defective and calls up your neighbors and starts fires in their toasters. And because parts of the toaster were built overseas in a country that was unfriendly to the United States, if you use the toaster in any government kitchen, it simply doesn't work. On top of that, the manual is badly written. Customers who buy it don't really understand how to use it. They attempt to toast jello, they use it in the bathtub. And when all of the various disasters occur, the fires and deaths, they find that the disclaimer that shipped with the toaster says that the vendor has no responsibility for anything that the user may do with the toaster, and therefore they have no legal recourse, and there is no penalty that comes back on the vendor.

However, the toaster does very well in the marketplace because it is cheaper than the other toasters that aren't certified and happen to work without fault. Those who go out and buy in large quantity, using the lowest bid, proceed to make you the market leader.

Certification does not guarantee that what you have is safe. It says that it meets the standards for the certification. It also does not tell you that it is going to be used safely or in an environment where it is appropriate to use it. That is why the Common Criteria is not appropriate.

Quality needs to be built in from the very beginning. As has already been noted by others, we don't know how to do that well, because this is an area that has been underfunded. It is an area where we need more research. We need more resources put into the agencies that are involved in this, particularly NIST.

This is a problem that we are going to continue to face for many years because we have such a large base of legacy code that is already in place and is too expensive to replace with something, even if we developed it tomorrow, that was much better.

Thank you for—the committee for listening to us on this today, and I stand ready to answer questions.

[The prepared statement of Mr. Spafford follows:]

Testimony before the House Government Reform Committee
Subcommittee on Technology, Information Policy,
Intergovernmental Relations and the Census

Exploring Common Criteria: Can it Ensure that the Federal
Government Gets Needed Security in Software?

17 September 2003

Statement of
Eugene H. Spafford
Professor and Director
Purdue University Center For Education and Research in Information Assurance
and Security (CERIAS)

Co-Chair of The U.S. Public Policy Committee
of The Association For Computing Machinery (USACM)

Member of the Board of Directors
of the Computing Research Association (CRA)

Table of Contents

Introduction	1
The Common Criteria: Pluses and Minuses	4
The Environment	7
The Human Factor	8
Some Recommendations	10
Conclusion	11
Acknowledgments	12

Introduction

Thank you Chairman Putnam and Ranking Member Clay for the opportunity to testify at this hearing. It is clear that there is a large and growing problem with the security of our cyberinfrastructure. Nowhere is this more apparent than in the computing systems used within the Federal government. Many of us working in information security have been alarmed for years by unfortunate trends in the way software is produced, acquired, deployed, and then used. High on the list of concerns has been the continuing poor quality of software, and in particular COTS (Commercial Off-the Shelf) software. This committee is to be commended for the series of hearings that it is holding on these issues.

This particular hearing poses the question "Exploring Common Criteria: Can it Ensure that the Federal Government Gets Needed Security in Software?" I will explain that the answer to that question is "no."

By way of introduction, I am a professor of Computer Sciences at Purdue University, a professor of Philosophy (courtesy appointment), a professor of Communication (courtesy appointment) and the Director of the Center for Education and Research in Information Assurance and Security. CERIAS is a campus-wide multidisciplinary Center, with a mission to explore important issues related to protecting computing and information resources. We conduct advanced research in several major areas, we educate students at every level, and we have an active community outreach program. CERIAS is the largest such center in the United States, and we have a series of affiliate university programs working with us in Illinois, Iowa, North Carolina, the District of Columbia, Ohio, Virginia, and New York State. CERIAS also has a close working relationship with over a dozen major commercial firms and government laboratories.

In addition to my role as an academic faculty member, I also serve on several boards of technical advisors, including those of Tripwire, Arxan, Microsoft, DigitalDoors, Unisys, and Open Channel Software; and I have served as an advisor to Federal law enforcement and defense agencies, including the FBI, the Air Force and the NSA. I am currently a member of the Air Force Scientific Advisory Board, and I have been nominated for membership on the President's Information Technology Advisory Committee. I have been working in information security issues for 25 years.

I began this document by listing my affiliations with ACM and CRA. This testimony is not an official statement by either organization, but is consistent with their overall goals and aims. ACM is a nonprofit educational and scientific computing society of about 75,000 computer scientists, educators, and other computer professionals committed to the open interchange of information concerning computing and related disciplines. USACM, of which I serve as the co-chair, acts as the focal point for ACM's interaction with the U.S. Congress and government organizations. USACM seeks to educate and assist policy-makers on legislative and regulatory matters of concern to the computing community. The Computing Research Association is an association of more than 180 North American academic departments of computer science and

computer engineering, industry and academic laboratories, and affiliated professional societies. The CRA is particularly interested in issues that affect the conduct of computing research in the USA. Both organizations stand ready to provide expertise and advice upon request.

Recent events, including the so-called Sapphire, Blaster, and SoBig worms have only served to underscore the vulnerability of our computing systems. Defacement of WWW pages and the spread of annoyance viruses pales in comparison to the potential for damage suggested by news that recent incidents may have affected some of our power systems during the August 15th blackout, that banking networks were unavailable because of contamination, and that sensitive military and law enforcement computers were also affected. As we increase our reliance on computing and networks, our potential vulnerability also increases. Deployment of new technologies, such as replacing some of our telephony with voice-over-IP (VoIP) suggests new reasons to be concerned about computer and network vulnerabilities.

Examination of many of the attacks, tools, and incidents that have come to light over the last few years indicate that far too many of the incidents have been the result of flaws in deployed software. What is even more discouraging is that so many of these flaws are likely the result of carelessness.

My research center (CERIAS) maintains a database of vulnerability reports and patches: The Cassandra Service.¹ On September 14, I performed a search of Cassandra. I selected the 100 software products with the largest number of reported vulnerabilities between January 1, 2000 and September 14. There were 2255 vulnerabilities reported for those 100 products in that time period — an average of over 11 per week. Overall, Cassandra has entries for over 5000 total vulnerability reports. Some of the programs had over 100 flaws reported in this period — an average of one every two weeks. Based on prior experience, there are undoubtedly many more flaws that have yet to be found, or that have been found and not yet reported to authorities

Analysis by my research staff earlier in the year revealed that over 20% of the reported flaws archived in Cassandra were caused by the failure to properly check for bounds on buffers, resulting in the possibility of attack via a “buffer overflow.” This is when an attacker provides more input than the program was coded to accept, and the programmer failed to properly account for the extra input. The result is that arbitrary data can be sent to overwrite program code or control information, leading to exploitation of the system. Analysis of data in NIST’s ICAT database² provides a similar figure, with 22% of the reported flaws being buffer overflows.

This is discouraging as buffer overflow is one of the first things we teach our students to avoid when they take their beginning programming classes. It is simple to avoid. It is a flaw that we have known about as a threat to both reliability and security for over 30 years. The infamous Morris Internet Worm of 1988 exploited a widely publicized buffer overflow, and, sadly, 15 years later we are still seeing commercial software written with buffer overflows. Vendors are

¹ This is named after the Trojan woman of legend who was cursed by the gods to have the power of prophecy but to never be believed. She warned the leaders of Troy not to take the wooden Trojan Horse inside the walls, but they did not listen to her. Troy was destroyed and Cassandra killed as a result. Cassandra access is free, and is available at <<https://cassandra.cerias.purdue.edu>>.

² See <<http://icat.nist.gov/icat.cfm?function=statistics>>

failing to learn from the past.

Further analysis on Cassandra data revealed that another 27% of the reported flaws were from other forms of failure to validate input (e.g., check bounds or correctness of values), and that other forms of simple design error accounted for an additional 26% of the flaws. Thus, nearly 3 out of every 4 reported flaws was the result of programmers making simple mistakes that have well-known causes and have been known and taught about for several decades.

It is clear that much of the software being used today in government and industry is severely flawed. What is more, the pressure of the marketplace has done little to eliminate these flaws. In some cases, it might be conjectured that the pressure of the marketplace has actually exacerbated the problem:

- Pressure by consumers for new features has led vendors to increase the complexity of software beyond the point where it is understood or completely testable. Companies with products having fewer features and reduced complexity are penalized as the population has sought out products with additional features even if they are not needed.
- Pressure for a reduced time-to-market has led to products being rushed in design and production, and then shipped before adequate testing has been performed. Companies with slower release cycles are viewed as “not innovative enough” and penalized in the marketplace.
- Efficiency of scale has resulted in vendors producing monolithic offerings that contain all features and options rather than specialized versions tailored for individual markets. It is not at all clear that the same software should be deployed in both environments. Consider that the same PC operating system, text processing system, spreadsheet, WWW browser and database system is likely to be found running on systems that contain and protect our nuclear secrets and intelligence information, and also on a home PC with someone's recipe book and baby pictures.
 - ◊ This results in reduced cost for the vendor, requiring only to produce, package, and document a smaller number of systems.
 - ◊ This results in reduced purchase cost (but not necessarily reduced operational cost³) for large customers as they can use the same hardware platform, add-on products and user training throughout the enterprise, despite wide variation in use of the product.
 - ◊ This increases the opportunities for an attacker because he can obtain and test attacks against the same software that is running in sensitive applications.
- Increased testing and better software engineering methods require hiring better-trained personnel, employing more expensive tools, and spending more to assure quality. End users have not shown a willingness to pay extra in support of this quality, thus putting companies at a disadvantage if they increase the care with which they develop code.

³ Many firms and government organizations separate operational costs from acquisition costs. Thus, a system that is \$100 cheaper per seat, but requires \$1000 more per year in after-market patching, anti-virus software and helpdesk support is often purchased as a means of “saving money.”

The lack of any meaningful penalty for producing flawed software (e.g., liability torts) has meant that there has been little in the way of normal market pressures to counteract the above.

The same factors may be indirectly affected academia's ability to train better software engineers. If a college or university were to commit to teaching its students more stringent software engineering practices and the use of more reliable programming languages, those students would not be as attractive in the marketplace: they use methods that are viewed as more expensive and less portable than current practice. Before long, the college would find it would no longer be attracting enough students to continue its program. Furthermore, corporate donations might well be reduced, increasing the cost of the program to maintain.

Similar pressures would — and do — impact research. Government and industry funding seems to be heavily oriented towards new and more efficient patching and protecting existing systems. This is not unexpected because so much has already been invested in those systems. However, that limits our ability as researchers to investigate new and perhaps more robust architectures and approaches, and thus serves to perpetuate some of the same flaws that are haunting us today.

The Common Criteria: Pluses and Minuses

Over the course of the last 30+ years there have been a series of standards and certifications for software quality assurance. DOD-STD 5200, the TCSEC, ITSEC, Federal Criteria, FIPS-140 and others have each attempted to address the need for strong protective measures in IT for government use. Each has had successes and flaws. The Common Criteria (CC) effort is the latest in this line of standards, and draws from the experience gained with those other efforts. It contains a number of much-needed enhancements over previous efforts, and has achieved widespread acceptance around the world.

For purposes of this testimony, I will assume that the Committee has had some form of tutorial on the Common Criteria and is familiar with its basic structure and terminology. As such, I will only address some major points about it here.

The Common Criteria is intended to provide a formal way to describe and evaluate the security properties of a software artifact. This has value when trying to understand the particular strengths of a product. However, the Common Criteria has a number of drawbacks that have been noted by those who have studied it. If we were to consider the impact of requiring Common Criteria evaluation for all software sold to the government, many of the disadvantages come to the fore.

- Malicious code hidden in software, such as might be introduced in offshore coding houses or by domestic criminals, is difficult for even trained auditors to find. Certification at any level below EAL-6 would probably not find such additions, and a clever programmer might still be able to hide some inside code that would evaluate at the highest level: EAL-7.
- Certification is oriented towards evaluation of products, not working systems. Thus, combining products together as occurs in normal operation may result in a system that is not certified, and may actually result in a system with weak overall security.

- Certification under CC is heavily weighted towards examination of documentation and vendor-produced artifacts. The lack of detailed, 3rd-party testing can be viewed as a significant weakness in the evaluation process.
- There are too few well-articulated protection profiles (PP) defined against which security targets (ST) can be written and products evaluated. To be useful, PPs should be written for specific environments and take into account a full risk assessment, and this is difficult to do. For this reason, PPs are difficult to develop and certify. PPs that are too general are not likely to be helpful. PPs describing unrealistic environments are similarly unhelpful.
- Certification for systems in high-risk environments, which are many of the government systems of most concern, should have certification at level EAL-5 and above. No certifications have been done at these levels, nor are the standards yet defined to support performance of such an evaluation.
- Although some systems might require evaluation at EAL-6, other systems would not require certification at such a high level because of their usage or risk level. Defining the appropriate level and protection profile needed for each application would be very time and labor intensive.
- Certification is currently expensive and time-consuming. Requiring CC certification of products might well keep them out of government use for years while the commercial sector is using the more recent versions.
- The government market is small enough in some sectors that vendors may choose to forego selling to it rather than undergo the time and expense of certification. This could prevent government agencies from taking advantage of new, more capable software and from interoperating with widely-used commercial software. This was often the case in the 1970s and 1980s under previous certification regimes.
- Software could be required to be certified at a higher level than necessary. For the vendors to comply, the cost of development and testing would need to be borne by the customers and end-users of that software, thus resulting in higher prices.
- There are not enough CC certification facilities to handle the load of products to be tested. Specialized training is required to produce qualified personnel to staff such facilities so there would likely be a severe backlog of products awaiting certification, exacerbating several of the other problems described in this list.
- The CC fails to consider standard product lifecycles. As such, products are often out-of-date by the time they are certified, and may not be the most mature, tested version of the software. The higher the evaluation level, the larger this lag becomes.
- Although there are mechanisms for certifying code that undergoes patching and after-market customization, it is not obvious if software with extensive patching needs and customization will be adequately covered under this mechanism. Current software does undergo frequent patching and upgrades.
- New, experimental and boutique products currently gain market traction by showing their

worth in controlled trials. If CC certification is required, the time and expense may keep these products from fair trial, and thus prevent them from succeeding in the market. This also includes university prototypes and experiments that may be deployed in new and novel manners. Many of the most important security tools in use in the market today would not be considered because they had their beginnings as small, leading-edge products.

- There is no obvious provision for open source software to be certified using the Common Criteria. The cost of certification and the required effort to produce an STare both prohibitive. Although open source is not better, per se, for security, there are some open source products that may be more appropriate for use in some circumstances. The current approach to CC certification would almost certainly exclude those products.
- Even certified code can support and spread viruses and other malware.
- Even certified code can contain flaws that can be exploited by attackers.

To the positive, however, the discipline brought to development to support eventual Common Criteria evaluation can be valuable. If more vendors actually did target risk assessments, thought about attacks and defenses, developed formal requirements documents, ensured that documentation was correct and up-to-date, used good production methods, etc, then our software infrastructure would be greatly improved. In fact, it is the *process* of software development rather than the *product* that may be most important in quality assurance. A quality process is more likely to result in a quality product. This fact is well-known in software engineering, and in engineering in general, but seems to be given little attention in many commercial software settings.

As an illustration of some of the issues, consider Microsoft's Windows 2000 product⁴. It was certified in October of 2002 at EAL-4+, which is currently the most rigorous certification available. However, this same system was the target of the Blaster worms, has been victimized by dozens of viruses, and has been the subject of several dozen patches for serious vulnerabilities discovered since the time of its certification. Although it is undoubtedly a safer product than it would have been without the effort to have it certified, it is certainly not what most people would consider a "secure" system. The value of the certification is therefore unclear.

Another aspect of this certification that should be noted is the protection profile (PP) against which Windows was tested. The PP defines the usage and threat model, and the choice of PP significantly impacts the interpretation of any certification results. Windows 2000 was certified against the CAPP protection profile. Dr. Jonathon Shapiro has posted an essay on his WWW site (<<http://eros.cs.jhu.edu/~shap/NT-EAL4.html>>) that addresses the certification of Windows 2000 at the CAPP/EAL-4 level. The following is a quote from his essay:

The Controlled Access Protection Profile (CAPP) standard document can be found at the Common Criteria website. Here is a description of the CAPP requirements taken from the document itself (from page 9):

The CAPP provides for a level of protection which is appropriate for an assumed non-hostile and well-managed user community requiring protection against threats of inadvertent or casual attempts to breach the

⁴ This is not meant to suggest that Microsoft's products are any better or worse than that of other vendors. This example was chosen to represent a product widely used within government and industry.

system security. The profile is not intended to be applicable to circumstances in which protection is required against determined attempts by hostile and well funded attackers to breach system security. The CAPP does not fully address the threats posed by malicious system development or administrative personnel.

Translating that into colloquial English:

Don't hook this to the internet, don't run email, don't install software unless you can 100% trust the developer, and if anybody who works for you turns out to be out to get you you are toast.

In fairness to Microsoft, CAPP is the most complete operating system protection profile that is presently standardized. This may be the best that Microsoft can do, but it is very important for you as a user to understand that these requirements are not good enough to make the system secure. It also needs to be acknowledged that commercial UNIX-based systems like Linux aren't any better (though they are more resistant to penetration).

Dr. Shapiro also notes that if he were to write a program that did nothing but paint the screen black (or blue), then it would be possible to have it certified at an EAL-4 or even EAL-7 level. All that the certification would mean is that the program reliably paints the screen, as designed and documented. There is really no conclusion that can be reached about the underlying fitness for purpose or utility of the software.

The Environment

The environment in which we currently deploy computing adds to many of our problems. We have a large base of legacy software and hardware that cannot be replaced quickly. In some cases, we would face incredible difficulty in replacing key systems, as we discovered during the search for Y2K problems. This means that any solution that applies to future systems yet to be produced will only gradually have an impact on the overall problem.

It is the variety of systems and operational environments that adds yet another layer of concern to security. We have such complex interactions and interconnections we cannot begin to understand all of their nuances. Years ago, my students and I described a series of security flaws that we labeled as "emergent faults" in our research. These come about when multiple systems, each operating correctly and securely according to their design, interact with each other and with the environment in unexpected ways to produce a failure. No single thing has gone wrong, but the combination has resulted in catastrophic failure. Even if several software artifacts are certified as Common Criteria compliant, once they are put together on a system or network it may be possible for them to interact in unexpected manners, leading to cascading problems.

We must also consider the effect of accidental physical failures. If there is a power failure in a government building, is it possible that a security monitor or firewall shuts down first in such a way as to leave other machines unprotected? Is it possible that a disk failure can result in an audit trail being lost? What of a lightning strike inducing current in a network line, a broken water main spilling water into a large data storage server, a failure of air conditioning resulting in overheating and failure of the main VoIP switch? In each case, the failure of the system is one that can lead to loss of data or processing and the presence of Common Criteria certified software does not prevent the damage.

There is also the case of physical failures from malicious actions. Arson and theft are two

malicious acts that can have serious impacts on sensitive computer systems. The theft of a diskdrive with personnel information, law enforcement information or building plans could all be catastrophic, but certification of software with Common Criteria would not necessarily mitigate the problem; if the underlying protection profile did not specify strong encryption for the stored data, there might well be little in the way of safeguards on the stolen data. An attacker could also cause the loss of power or water line break described above, either as a destructive act, or to disable a particular set of mechanisms to enable some other activity. The CC will not provide protection against such acts.

The Human Factor

It is a basic fact that we would not have computer security problems if there were no people. Although that may sound flippant, the point is that human beings are the ones who use our IT systems, and humans are the ones who abuse them. We need to pursue approaches that reduce the risk from both approaches.

Humans abuse computer systems by breaking into them, committing fraud, causing denial of service, writing viruses and other malware, committing identity theft, and committing a long list of other criminal activities. No matter how good the technology may be, there will likely always be ways to abuse it. If nothing else, "insiders" who have access to the computer systems are in a position to misuse their authority to access or change sensitive data. We have many historical accounts of traitors accessing national defense information for foreign powers, of law enforcement agents accessing details of confidential informants and undercover operations in return for bribes, and of government officials accessing personal data for inappropriate uses.

Technology can only provide protection for information up to a point. Once the information is accessed by people and other applications with appropriate authorization, it is necessary for non-software methods of protection to be applied. Thus, personnel security mechanisms must be in place to appropriately screen individuals before they are placed in positions of trust, and to periodically reevaluate them. Operational security (OPSEC) methodologies should be applied in the maintenance and operation of the systems to audit usage, detect questionable behavior, and respond appropriately.

One of the most important aspects of information protection is law enforcement. There needs to be a credible threat of discovery and prosecution in place to deter individuals who are considering misuse of IT systems, and to appropriately punish those individuals who do commit transgressions. This requires trained investigators, adequately-equipped laboratories, and the necessary personnel and financial resources to pursue investigation and prosecution.

The current state of law enforcement for cybercrime is far below the level needed to provide an adequate deterrent to criminal behavior. Although there are a number of well-trained investigators at the Federal level, their numbers are far too few to deal with the many cases brought to their attention...or the many more that would be brought to their attention if the victims had some hope of the cases being pursued. There are also a few well-equipped forensic laboratories in the

US, but they are few in number and heavily loaded with casework. There are no nationally accepted standards of training for those examiners, there are a limited number of software tools to use in investigation, and there are few resources being expended to advanced research and training in the area. The cost and complexity of investigation and prosecution is such that officials are often reluctant to pursue cases without proof of a large, obvious loss.

The situation at the state level is even more dismal, with only a few states possessing some advanced resources that can be applied to computer crime investigation. When one considers that many Federal crimes are first discovered as crimes against local entities, and that many offenders are juveniles who are unlikely to be prosecuted at the Federal level, the need for adequate state resources becomes more acute.

The other aspect of human behavior that leads to IT system compromise is that of non-malicious activity brought about by stress, carelessness and/or ignorance. Humans who operate the IT systems may not have sufficient training or resources (particularly, time) to appropriately operate the systems in a safe manner. This becomes a major issue when every physical desktop holds a high-end, networked computer running a complex and (probably) flawed operating system. Small errors in configuration of a single networked system can result in catastrophic, cascading failures of other systems in the enterprise. The ease with which new software can be downloaded and executed (either intentionally, as a browser plug-in, for example) or unintentionally (as occurs with many email viruses) is one source of many security problems. A system certified to a high level in the Common Criteria can still fall prey to these problems unless they are all anticipated and addressed in the design (and in the security target, ST).

Currently, secretaries and low-level administrative personnel are typically equipped with high-end computers running full operating systems capable of spreading viruses and denial of service attacks, when all they may really need is access to a mail program, address book, and WWW browser. The mix of extra functionality combined with the lack of training provides a potent — and dangerous — combination. A more accessible solution than trying to ensure the quality and safety of all the software running on every platform is to explore systems with more limited functionality, including use of “thin-client” systems that remove the main computing platform and unnecessary utilities from the reach of the inexperienced user.⁵

It is also the case that in many environments the administrators of systems are charged with the security of the systems they administer, but they are not given training or support for those tasks. The result is that they are under time pressure to install fixes, configure security tools, follow up on incidents, and respond to user concerns in addition to their regular duties. Without adequate training and tools, they cannot respond effectively nor in a timely manner. Although well-intentioned, they end up letting some critical tasks (such as patch application) slip because of a lack of time.

Coupled with all of this is the overall problem of user interface. Too many products are

⁵ A thin-client is basically a terminal or limited workstation connected to a larger computer system. This architecture also brings other benefits, including great ease of administration, more complete archiving of software, easier patching, better control over software licensing, and greater resistance to viruses.

provided without easy-to-understand (and easy-to-operate) controls for the protection mechanisms that are in place. To enable a built-in firewall, for instance, might require as many as a dozen operations to set various options and variables. Not only must those settings be found, but they must be understood — what is the indirect effect of each setting, and what else needs to be configured? Too often the documentation is unhelpful, incomplete, or incorrect. Only some of this is addressed by the Common Criteria process, and even then there is no formal process of determining a match of the user's abilities with the user interface of the security mechanisms. It is little wonder that so many computing systems are configured incorrectly and dangerously.

Some Recommendations

There are several actions that can be taken to reduce the threat of abuse of government computers. All of these can be derived by examining the problems that confront us. These are independent of the Common Criteria. Among those that have the highest likelihood of making a difference, I would include:⁶

1. Emphasize the need for a systems-level view of information security. Assuring individual components does little to assure overall implementation and use. This requires trained personnel with an understanding of the "big picture" of IT security. Too often those who design and specify the systems do not understand how they are actually used....or misused.
2. Establish research into methods of better, more affordable software engineering, and how to build reliable systems from untrusted components. 15-20 years ago the decision was made to cede research in this arena to the commercial sector, believing the market would drive innovation. That has not happened.
3. Increase the priority and funding for basic scientific research into issues of security and protection of software. Too much money is being spent on applying patches to intrinsically unsound systems and not enough is being spent on fundamental research by qualified personnel. There are too few researchers in the country who understand the issues of information security, and too many of them are unable to find funding to support fundamental research. This is the case at our military research labs, commercial labs, and at our university research centers.
4. Explicitly seek to deploy heterogeneous environments so that common avenues of attack are not present. This *may* require some extra expense *at first*, but eventually it may lead to increased compliance with standards, increased innovation, and increased choice in the marketplace, thus lowering costs while increasing security. If real standards (rather than de facto standards) are developed and followed, interoperability should not be a concern.
5. Complementary to the previous recommendation is giving thought to different architectures. Rather than a computer on each desktop, thin-client technologies based on a mid-size computer in a centralized location can provide all the same mission-critical

⁶ I provided a similar list to the House Armed Services Committee Subcommittee on Terrorism, Unconventional Threats and Capabilities in my written testimony of 24 July 2003. Although this document does not address all of these points, I believe they are still worth considering.

services, but remove many of the dangerous aspects of distributed PCs. For instance, patches need only be applied in one location, and there is a greatly reduced possibility of untrained users loading untested media or software.

6. Rethink the need to have all systems connected to the network. Standalone systems may not receive all of the latest patches as soon as they come out. However, that alacrity may not be needed as those systems can no longer be attacked over the network.
7. Require greater efforts to educate personnel on the dangers of using unauthorized code, or of changing the settings on the computers they use. It is still often the case that personnel will turn off security features because they feel it slows them down or gets in their way. Unfortunately, this can lead to significant vulnerabilities.
8. Revisit laws, such as the DMCA, that criminalize technology instead of behavior. It is extremely counterproductive in the long run to prohibit the technologists and educators from building tools and studying threats when the “bad guys” will not feel compelled to respect such prohibitions.
9. Provide increased support to law enforcement for tools to track malware, and to support the investigation and prosecution of those who write malicious software and attack systems. This includes support for additional R&D for forensic tools and technologies.
10. Do not be fooled by the “open source is more secure” advocates. Whether source is open or proprietary is not what makes software reliable. Rather, it is the care used to design and build it, the tools used to construct and test it, and the education of the people deploying it. In fact, some Linux distributions have had more security flaws announced for them in the last 18 months than several proprietary systems. However, some open source software, such as OpenBSD and Apache, appear to be far more reliable than most proprietary counterparts. There is no silver bullet for problems of quality and security, and that includes the Common Criteria.
11. Initiate research into the development of metrics for security and risk. Acquiring systems based on cost as the primary criterion is not reasonable for mission-critical applications. We need to be able to differentiate among different vendor solutions, and set standards of performance. Common Criteria evaluation is not sufficient for this purpose, especially when systems are evaluated against very different protection profiles.
12. Establish better incentives for security. The current climate in many government agencies is to penalize operators for flaws, thus leading many of them to dread enhancement and exploration of better security.

Conclusion

It is clear that we have deficiencies in our cyber defenses. Poorly designed and incorrect software poses a particular threat because it can be so widely deployed in government and the civilian sector. We need to find better ways of increasing the quality of the systems we purchase and deploy. For the reasons given in this testimony, application of the Common Criteria cannot ensure that software used by the Federal government will provide a sufficiently secure

Mr. PUTNAM. Sounds like this could be a fun panel. Mr. Clay will lead off this round of questions.

Mr. CLAY. Thank you, Mr. Chairman. And I will start with Mr. Spafford. You said your research center has a data base of computer vulnerabilities. In your search of that data base, how many of the software products identified were certified under the Common Criteria?

Mr. SPAFFORD. Sir, I didn't do a search specifically on that criteria, but from the numbers that I know from looking at similar searches, I would suspect that we have several hundred that apply to certified products. As I noted, there are over 100 for Windows. A few of the other products, the firewalls and intrusion detection systems, the Oracle data base system has a few as well. We know that there are several hundred vulnerabilities for the bulk of certified products.

Mr. CLAY. You also point out that software certified at the highest level of the Common Criteria is subject to the same worms and viruses as software that has not been certified. The certification process is already a long and costly process. What would have to be changed in the Common Criteria to address these problems, and what would that do to the cost and time for certification?

Mr. SPAFFORD. Sir, the certification process is against the documentation that is provided by the vendors and against a set of specifications that have been put out by the groups that have set the Common Criteria, and those do not include issues such as resistance to malicious software.

There are architectural issues to the way that the code is actually written that would need to be changed in the products fundamentally. So, for instance, taking macros out of word processing and spreadsheets, preventing e-mail programs from automatically executing attachments are ways to stop viruses, but they are not the only ways to stop those kinds of software problems.

And those are not issues that are tested under the Common Criteria. Those are architectural features that are actually part of the product and the reason it is sold as it is.

Mr. CLAY. Mr. Klaus, you indicate that you do not believe that Common Criteria evaluation improves security. How would you propose the government determine in the procurement process that the software it buys is secure?

Mr. KLAUS. I think the—in talking about the security products, say, for example, the firewalls and the intrusion detection systems, today there is no certification that I am aware of that says, you know, take Gene Spafford's data base of thousands of vulnerabilities and exploits and different ways hackers get in today and evaluate whether a firewall or IDS system stops all of these vulnerabilities and all of these attacks.

So today there is no benchmark or operating system among the security products to say which one is most robust against these types of attacks; these are the known attacks much less the unknown attacks that are continually evolving.

If we can just measure how good are security vendors keeping up with the current pace of vulnerabilities, because there is a lot. There are over 200 vulnerabilities, like I said, every month, where

we are tracking and need to keep measuring the quality of the security vendors' products.

It is a little bit counterintuitive. I think if we look at some of the commercial certification companies out there, they have been able to hit the goals. When you look at the companies that certify the antivirus companies, they meet 99.9 percent of all viruses. They have been able to hit it. So they are testing what is out in the wild, what are the latest things that are happening, so they can quickly, at the end of the product, measure did the security company keep in pace at the very end? Did they hit the end result of what they said that they would do for protecting against those threats?

And then the onus of having a very robust security product and the processes are still left on the security vendor to follow. And from a speed and cost perspective, because it is only testing at the very end, did this thing catch all of the hacker exploits, all of the worm exploits, all of the different ways that systems get compromised, it is a much more lightweight process. You can accomplish it in a month. It doesn't take a year to go through that, and therefore the turnaround is much faster.

Rather than trying to be completely overcomprehensive in your evaluation of every detail and aspect of the security design and architecture, I think that needs to be held to probably the security vendor themselves making sure that they end up with that, because otherwise they themselves become part of Gene Spafford's data base of vulnerable systems.

But, on the flip side, the most important thing in terms of protecting the government: Can they stop these risks? I think shrinking it down to a much more focused process would help drive lower costs, faster speed and a much more accurate measurement of is this product more secure or less secure for the government.

Mr. CLAY. Ms. Davidson, did you have something to add?

Ms. DAVIDSON. I did. Actually I had a couple of responses to that, one of them a personal anecdote. My company did look at deploying one of the hottest new security products in a particular sector, which is supposed to defend against certain classes of application vulnerabilities. This is something, a specialty firewall, that you would put out to protect yourself against various types of attacks. It claimed to be the one of the market-leading products. My hacking team broke it in less than an hour using an attack that product was supposed to prevent.

It is important that security products, because they are the early warning system, have some type of independent assessment of security worthiness.

I am also aware that people's intrusion detection systems failed when Slammer was going around because of the composability of the systems. They were running things back end which themselves were not secure. I would certainly be open to flexible ways of validating the security worthiness of security products, but it is not all about feature function. It does one no good to protect, allegedly protect, against certain classes of attacks only to find that the security system itself is badly flawed, and that in at least two cases has been our experience.

Mr. CLAY. Thank you.

Mr. Thompson, did you want to add something?

Mr. THOMPSON. I just wanted to point out that the security evaluation process is designed to verify what a vendor claims, but it does that—there is a very publicly available statement of what the vendor claims. For example, if—you know, if the toaster manufacturer says that—has this evaluated as a bread storage device, it would be evaluated as a bread storage device. And if the government wanted to buy toasters, and they wrote a protection profile that specified toasters, the bread storage device probably wouldn't meet the protection profile for toasters. And somebody could write a security target for a bread storage device, and it would be—you know, it would be classified as a bread storage device.

The confusion—the CC process allows products to be compared by specifying their security criteria in a semiformal language that is easily comparable.

Mr. CLAY. Well, the security issue sounds more like a moving target, you know, as people come up every day with new viruses, new worms, new ways to penetrate computers.

Mr. THOMPSON. Well, that is sort of a—

Mr. CLAY. Can we win? Can we win the battle of securing these computers?

Mr. THOMPSON. That is a different—finding patches and fixing them is a very difficult process, very expensive process, and I don't think that is the way we are going to win in the end. There is certainly things we can do—to find a patch is something we should do, as long as we have these vulnerabilities, but the kind of software development that the Common Criteria is encouraging is using sound engineering principles and design life cycle processes.

And that is with the higher assurances like EL6 and 7, encourage those kinds of things. In other words, you can't—if you are going to evaluate EL7, you have to develop it in the process of documenting. You have to formally prove that it meets its security policies and things like that. And those engineering principles have to be applied to the development process. And we think that is a more—in the long run the only way you are going to get secure products.

Mr. CLAY. Mr. Klaus.

Mr. KLAUS. I think there is—within the Common Criteria, one of the issues that I think that the previous panel had really pointed out repeatedly was that this is still an art form in terms of finding the vulnerabilities, more R&D money for automating the tools.

But at the end of the day, what we are finding is, is it really requires subject matter experts to be able to—who understand how to find buffer overflows, how to find heap overflows, how to find—many of the techniques that hackers use to break into the systems.

And what we find is, a lot of the approved testing labs don't have that expertise, to find these kinds of vulnerabilities; and from that perspective, we are not measuring whether the systems have those types of vulnerabilities. And I think if we can build a system that's measuring for “can the security products find and identify attacks and stop them”—I think that's where Mary Anne Davidson was pointing out that security products need to get better at: on “being evaluated,” on “can they identify the attacks?”

Just as important as identifying the attacks, it is almost more important for enterprises to make sure that we're also not identify-

ing false positives. This is where we falsely see, say, legitimate traffic and identify it as, “oh, here’s an attack,” but the minute you do that, you start cutting off real business transactions and so on, and then your security product is no longer trusted; or you turn off that functionality within the security product, and you are now less secure.

And the other thing that needs to be tested is for invasion techniques. A lot of the known hacker community has published—there’s lots of white papers on how to evade many of the security products, and many of the security product vendors behind that have not responded to those techniques. They are still valid and still work.

And there’s no, I guess, in the certification process, anything that reflects how good they identify the attacks, the false positives, the invasion techniques; and I think, to answer another question, I think would be critical for security companies to be measured on is the effect of “zero-day” exploits.

What I mean by zero-day is, when the worm comes out or a virus comes out, the most impact, full-time, is within the first 24 hours, in that it’s spreading and nobody has protection. All of the security vendors are trying to respond to get the latest, “What is that attack; OK, let’s update our security products.”

There’s a concept of—within the security industry that we’re moving toward behavior-based security models, where I don’t have to take a fingerprint of every virus, every worm. I’m actually looking at the behavior of that program so that if something tries to compromise a system and acts to propagate and format your hard drive and change your registries and other things on the system, those are all bad behaviors, and it gets flagged; and you could stop it without even knowing the—what was the virus before you saw it.

And I think if we added some measurement to how good do security products deal with zero-day threats, all you have to do is test an old version—if it hasn’t been updated, test against a new threat. Did it stop it? If it did, great, you get a point for that. If it didn’t, you don’t get a point and you can start measuring across a lot of security parts out there.

Ms. DAVIDSON. With all due respect, I think most of us believe in defense and depth and that security cannot be outsourced. If a vendor has a fault in their product, they cannot outsource the remedy for that, even to intrusion prevention.

For example, the customer comes to me and says they found a fault in our software. I can’t say to them, Do you have a fire wall? Do you have an intrusion prevention system? Because if you do, I won’t fix it. They will have my head.

So I have to get it right the first time anyway. And if I get it wrong, it will still cost me a million dollars to fix it if it’s on every single version of product on every operating system.

Everyone needs to write better code. In order to write better code, we need better tools. It’s not just training, because developers are human; they make mistakes. One mistake and the hacker is in.

Mr. SPAFFORD. I wanted to add, we have mentioned that we need more research and tools. We need more personnel. This is an area where we have a very small pool of expertise. But one thing that

would make a difference, I believe, is a matter of accountability. And the sentiment expressed by Ms. Davidson here has not been widespread enough within the industry, which is, if there is a problem in the code, then the people who wrote the code are held responsible.

Currently, when the government buys systems, if they have a failure, then everybody rushes around and applies a patch and then goes on as if nothing else had happened until the next failure and the next patch.

I really believe that if there's some negative feedback to the vendors involved, if they have a bad history of producing software that isn't reliable, then perhaps that should be figured into the next series of acquisitions. Perhaps there should be a penalty applied to some vendors if they consistently provide bad software. It's something worth considering because simply encouraging them by buying the next product cycle isn't resulting in the changes that we should be seeing.

We are seeing vulnerabilities that have been known for 30 years to be security problems and bad practice; and we are discovering that 50 percent of all the vulnerabilities that are being reported today, 2 or 3 a week, are those bad practices that are 30 years old and that my colleagues and I teach in the very first few weeks for students to avoid. There should be something back at—a negative pressure to have them start paying attention to better practice.

Mr. CLAY. Thank the panel for their responses.

Mr. PUTNAM. Could you give us an example of a 30-year-old vulnerability?

Mr. SPAFFORD. In the very introductory programming classes we teach, we tell the students they should check the inputs. For instance, if it's requested that a number be provided between 1 and 10, we ask them to check that the value is between 1 and 10. If they are asked to provide a character string that is 10 characters long, then they should check to make sure they aren't provided with one that is 11 characters or 1,000 characters.

When we talk about buffer overflows or when you've heard that mentioned by the panelists, that's a case where a program was expecting 20 characters and was given 2,000 and there was no check made to see that too many characters were provided. That is something that has been known for 30 years to be a problem. It has been exploited in many systems. We teach against it, and it's still occurring and being discovered at the rate of several a month.

Mr. PUTNAM. Ms. Davidson, Oracle began certifying products very early in 1998. What led you to come to that conclusion and how has it affected your business?

Ms. DAVIDSON. We initially began doing evaluations, actually the pre-Common Criteria days, we did the orange book and IT section evaluations. Actually, we did four of them at once on two of our products. We did that because one of our core customer constituencies demanded it; at least we thought they demanded it, as I testified previously.

They would occasionally wuss on the procurement requirement—that's a technical term—but we kept doing them anyway. We thought it was important. We found the benefits for us were substantial for the reasons that I previously laid out.

I feel actually the cultural values—making security part of a corporate culture has been the biggest value. I don't have discussions or arguments about, are we going to hold the release, because there's a security fault. Of course, we do that. This is something that we are measuring on and it is something we are held accountable on.

We certainly are not perfect. We have developers who have committed the sin of buffer overflows or not checking input conditions, but I would consider myself to be successful if I could stomp buffer overflows in our time, but we need help to do this. We spend a lot of money training people.

When I was in the Navy, there was an expression, "To err is human; to forgive is not Navy policy." People do make mistakes in programming. If there are 21 conditions that have to be validated, our developer checks 20 of them, the hacker only needs to find that one.

If complexity is the enemy of security, so are manual processes. The more that we can automate some of these checks in addition to training people and holding them accountable, the easier it will be for people to do the right thing. And right now it is really hard, because you are only as good as every single person checking every single possible programming condition; and they're not perfect, and they never will be perfect. It's been good for us as a business. And I don't think the Common Criteria is a solution for all security ills, but I think if people don't bake security into their development processes, however we get there—just checking air conditioners will also not make us secure; you need both.

Mr. PUTNAM. How would you respond to the toaster metaphor? You know your company is committed to it, you follow through, you are believers in Common Criteria. Mr. Spafford and to a certain degree Mr. Klaus have laid out a series of arguments why it will not get us where we hope that it will. How do you respond to that?

Ms. DAVIDSON. I think it is a great analogy on a lot of levels.

The other counter-argument is, if you glue all the pieces together, you may not get a secure house, but if you don't start off with a secure foundation, you certainly will not get a secure house.

Yes, evaluations do not make for perfect conditions. But would you really want to plug your toaster in, even with the cord, and have no idea how strong the building foundations were, whether the engineers had done their jobs, whether they had building inspectors in. You need to do a lot of things to have secure software.

I think evaluations are part of the answer because it will change the way people build software. It changed the way we build software. I think have better validation. If we had automated tools—most security faults are not faults in the security mechanisms; they are as a result of bad programming. If you have better automated checks for good programming practice, you also will be able to add a level of robustness.

And the third piece I mentioned earlier is, many vendors, despite our best efforts, don't deliver products that are secure enough out of the box. We give our customers long lists of things to do and to tweak to become secure. And most system administrators never have enough hours in the day to do that.

You have to make it easy for people. You have to install your product so that ideally people don't have to do anything—like the Cuisinart—to have it operate securely. If you do that, it not only lowers people-cost-of-operation and increases their security, I think you will get resistance to some viruses and worms that typically exploit lots of things that are left wide open on your system, or things that are left lying around in your system because a vendor shipped it and the customer didn't know how to secure it.

Common Criteria is a strong necessity, but I would agree it is not sufficient.

Mr. KLAUS. I think from the house perspective, if you look at how—I just went through the process of finishing a house. The certification process and compliance is typically at the end of the process. You have the government come out and look at the house and make sure it's up to code and you meet that criteria, or for a building or for—you are looking at the, at the very end, did the House meet all the necessary standards. And some of the important—and it's looking for the critical issues. Are sprinklers in place, etc.

What you don't see in the certification process—and this is where I think we are failing—is, the opposite is happening in the Common Criteria where if you had a document—as the architect, the designer of the house had to sit there and as it goes through the process add another year—I mean, it took a long enough time to finish the house. If you add another year to building the house, to make sure that everything was documented, and here was all—I trust my architect to make sure it's designed to be built strongly.

The government shouldn't have to go in there and say, did you use all the metrics to make sure it's going to stand, and then at the end the government checks to make sure the most important issues are addressed and certified. And to the extent—if we could move the Common Criteria more to, did the important issues get addressed?

And I think you could look at it, hey, a lot of these applications, especially business applications, are very complex. Many ways—many lines of code, etc. But if you actually identify what are the most common ways that hackers, worms, viruses hack into a system, the majority of the risk is at the network protocol code. You know, if you look at, why did Blaster get into the operating system, well, it was because there was an RPC service running on every Windows box that had this vulnerability.

You can say there's millions and millions and millions of lines of code within operating systems and these business applications, but the most important thing is to look at, what are the things that are exposed at the network level? That tremendously reduces what you have to evaluate.

We do a lot of security penetration tests, a lot of security assessments trying to figure out how would a hacker break into a system, and we always start at the network layer. And I think if the certification process looked more at—the same way that the hackers, the worms and viruses looked at, how does somebody break into the system, you'd start saying, OK, do you want to check the doors and the windows? You don't want to—I mean, you don't try to evaluate every wall and floor, the whole house. You evaluate the areas that hackers get into.

And that's kind of intuitive, but if we focused on the bigger issues measuring whether you have a good security product or not; less on, did the overall process get followed, because right now it's not helping us find the buffer overflows and other things within the product at the end of this certification process.

Mr. PUTNAM. Dr. Spafford, you started this metaphor, and I would ask for you to talk a little bit about what the better alternative is. Software assurance, how do we get there if it's not Common Criteria?

Mr. SPAFFORD. I didn't mean my comments to mean that Common Criteria is not a value, because I believe it is. It provides guidance as to how go about building a quality product. But it's building that quality product that is the key to what we are talking about.

It's not simply a matter of security. We want to have greater trust in our systems, but we also want it to be reliable in the face of failure, unexpected circumstances.

It appears, for instance, that the blackout that occurred in the East Coast was as a result of unfortunate circumstances happening at once, without sufficient capacity and reserve to make up for the failure. We don't want that to occur with our computer systems either.

That means going back and looking at fundamental assumptions that are made on how we build the systems. What are the features that we really want? How is it being built? Is it being built using good tools and by people who understand the technology? Are they putting in more features than are really necessary, which I believe is the root cause of a number of the problems that we see. Is the documentation in the interface? Are those two items put together in such a way that the average user is able to understand how to use the system and how to configure it?

Again, I do not believe that is the case. The average user currently is very often someone at home who doesn't understand what a firewall is or a virus is or what it means to have their system connected all the time.

Then we have to have better testing tools and some known reasons to test, some known test sets to work against. We have to be able to test in real environments, so that if we are going to deploy something in a large-scale system, we have to have testbeds to do that; and again, we have to have the people trained to do that.

And last of all, we have to have a mechanism so that we understand if we need to apply the technology to new arenas, how we go about going back in the process and changing the technology rather than simply reusing the old technology because that's what we have a large investment in. We should be using the most appropriate tools for the tasks at hand.

What has happened over the last 30 years for computing, if we look, there's been incredible strides from mainframes and small networks to where we are now with global, international activity with our systems. We don't even know where some of our software comes from because of the international trade and development that goes on.

We spent those 30 years trying to make the technology work, and I think we have done a really admirable job of that. So much of

our society, so much of our dominance in the world has come about through our ability to create good technology. But now we have to change our mind-set to think about how to most appropriately use that technology and how to make it safe, and that means really taking a step forward, leaving behind some of the technologies of the past.

So again, to summarize, it's not a simple step. It's going to be a whole number of steps throughout the life cycle of building and designing software. And to revisit Mr. Klaus' comments about the architect, well, the architect has been through many years of professional training. They probably served an apprenticeship with a master architect to understand what they're doing. And if they have designed his house, and it ends up getting built and there's no doors and it collapses afterwards, he has some recourse. And it's possible that architect will not be able to sell a design again in the future.

We haven't done that in the software arena. We need to start thinking in terms of how we're going to protect our future. Are we going to continue to reward bad performance?

So it's a long answer—I apologize—but it's a very multifaceted problem.

Mr. PUTNAM. The \$64,000 question: Should we expand Common Criteria to civilian agencies?

Mr. SPAFFORD. I believe that, on balance, that should not be mandatory. As a voluntary step, it may be good, but mandatory, it will not solve the basic problems. There are certified products that won't work as required.

The process is not easy to understand. The Common Criteria standard document is 700 pages long, and so many of the people are going to be buying and deploying these systems who won't understand what the certification means, which is why I used the analogy of the toaster. The average consumer won't understand what that means.

It can help to get some vendors to pay more attention, but I believe that the additional overhead, time and costs that were discussed by the other panelists are probably counterproductive to government's needs. I believe there are other steps that should be taken first.

Mr. KLAUS. My answer would be, until the, at least the 3 years we talked about, better measurement of what you are trying to do with the Common Criteria is met, meaning, is this product actually providing better protection against the known threats and making process later, because if it takes a year to get our products out the door to help the government, they're going to be a year behind the commercial sector. And the cost of it is—I think overall, the cost is expensive, so startups have—will have a hard time entering into the government sector.

But most importantly, if the cost was moving toward making the products better, I'd be in favor of it. Today there's very little value in what it is today.

Mr. PUTNAM. Ms. Davidson.

Ms. DAVIDSON. If it's not too expensive and doesn't take too long to do.

With all due respect to Mr. Klaus' company and their fine products, we have more complex products. We get certificates out within 6 months of the production release, and we release major versions of product every year to 18 months. It is cheap compared with the alternative.

We are already paying for bad security. I believe that it should be extended at least—clearly, to entities who have a national security focus. And if the Department of Homeland Security is not doing national security, what is it that they're doing?

As I mentioned earlier, there are other things we can do, but unless we fundamentally, as an industry, change the way that we build product, nothing will ever change. And this is the government's last chance on this. If we abandon information assurance efforts and go only to a testing approach, you will never know whether someone developed good product.

Testing alone, while I think an important add-on on top of Common Criteria evaluations, also will not solve the problem.

Mr. THOMPSON. I think I agree with Gene that we have to encourage good software development, good software design, and encourage companies to develop products that are safe in the beginning; and not just throw them on the market and let the rest of the world find the bugs one at a time or let the hackers find the bugs. They need to produce good software and need to be held accountable when they don't.

And the Common Criteria approach to evaluation encourages vendors to do that. It is not designed to find bugs in a particular release of a product, but designed to encourage vendors to use good software and build that house according to good architectural principles in the beginning and not to find, you know, where the beams have been left out or where the small beams were used.

Expanding the market for evaluated products would encourage that, send a signal to industry that the government is serious about good software engineering and development; and the products should be, you know, secure from the get-go. And anything you can do to allow—make vendors accountable to their—for putting out bad software would further the government's ability to buy good software. Everyone would have better software available if vendors were held accountable.

Mr. PUTNAM. Thank you, Mr. Thompson.

I want to thank all of our witnesses today, particularly the second panel, for their efforts in helping us to better understand this very complicated issue.

Gaining assurance that the software the government buys to protect itself actually can do the job is an important goal.

I also want to thank Mr. Clay and Ms. Watson for their participation. In the event that there may be additional questions that we did not have time for, the record shall remain open for 2 weeks for submitted questions and answers.

With that, the subcommittee stands adjourned.

[Whereupon, at 12:15 p.m., the subcommittee was adjourned.]

